



Switch Engine™ Command Reference Guide

for Version 32.3

9037676-00 Rev AA
December 2022



Copyright © 2022 Extreme Networks, Inc. All rights reserved.

Legal Notice

Extreme Networks, Inc. reserves the right to make changes in specifications and other information contained in this document and its website without prior notice. The reader should in all cases consult representatives of Extreme Networks to determine whether any such changes have been made.

The hardware, firmware, software or any specifications described or referred to in this document are subject to change without notice.

Trademarks

Extreme Networks and the Extreme Networks logo are trademarks or registered trademarks of Extreme Networks, Inc. in the United States and/or other countries.

All other names (including any product names) mentioned in this document are the property of their respective owners and may be trademarks or registered trademarks of their respective companies/owners.

For additional information on Extreme Networks trademarks, see: www.extremenetworks.com/company/legal/trademarks

Open Source Declarations

Some software files have been licensed under certain open source or third-party licenses.

End-user license agreements and open source declarations can be found at: <https://www.extremenetworks.com/support/policies/open-source-declaration/>



Introduction to the Switch Engine™ Command Reference

[Conventions](#) on page 3

[Related Publications](#) on page 3

[Send Feedback](#) on page 4

[Help and Support](#) on page 4

This guide is intended for use by network administrators who are responsible for installing and setting up network equipment. In addition to comprehensive conceptual information about each feature of our software, you will also find detailed configuration material, helpful examples, and troubleshooting information. Also included are supported platforms and recommended best practices for optimal software performance.



Important

Release 31.6 introduced new names for the network operating systems running on Universal hardware. ExtremeXOS (EXOS) was renamed to Switch Engine and VSP Operating System Software (VOSS) was renamed to Fabric Engine. All references to ExtremeXOS apply to Switch Engine 31.6 and later.



Note

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

Conventions

To help you better understand the information presented in this guide, the following topics describe the formatting conventions used for notes, text, and other elements.

Related Publications

ExtremeXOS and Switch Engine Publications

- [Switch Engine 32.3 Command Reference Guide](#)
- [Switch Engine 32.3 Feature License Requirements](#)
- [Switch Engine 32.3 User Guide](#)
- [Switch Engine 32.3 Release Notes](#)
- [ExtremeXOS Quick Guide](#)

- [Extreme Hardware/Software Compatibility and Recommendation Matrices](#)
- [Extreme Optics Compatibility](#)
- [Switch Configuration with Chalet for ExtremeXOS 21.x and Later](#)
- [ACL Solutions Guide](#)
- [Using AVB with Extreme Switches](#)

Open Source Declarations

Some software files have been licensed under certain open source licenses. More information is available at: www.extremenetworks.com/support/policies/open-source-declaration/.

Send Feedback

The Information Development team at Extreme Networks has made every effort to ensure that this document is accurate, complete, and easy to use. We strive to improve our documentation to help you in your work, so we want to hear from you. We welcome all feedback, but we especially want to know about:

- Content errors, or confusing or conflicting information.
- Improvements that would help you find relevant information.
- Broken links or usability issues.

To send feedback, do either of the following:

- Access the feedback form at <https://www.extremenetworks.com/documentation-feedback/>.
- Email us at documentation@extremenetworks.com.

Provide the publication title, part number, and as much detail as possible, including the topic heading and page number if applicable, as well as your suggestions for improvement.

Help and Support

If you require assistance, contact Extreme Networks using one of the following methods:

Extreme Portal

Search the GTAC (Global Technical Assistance Center) knowledge base; manage support cases and service contracts; download software; and obtain product licensing, training, and certifications.

The Hub

A forum for Extreme Networks customers to connect with one another, answer questions, and share ideas and feedback. This community is monitored by Extreme Networks employees, but is not intended to replace specific guidance from GTAC.

Call GTAC

For immediate support: (800) 998 2408 (toll-free in U.S. and Canada) or 1 (408) 579 2826. For the support phone number in your country, visit: www.extremenetworks.com/support/contact

Before contacting Extreme Networks for technical support, have the following information ready:

- Your Extreme Networks service contract number, or serial numbers for all involved Extreme Networks products
- A description of the failure
- A description of any actions already taken to resolve the problem
- A description of your network environment (such as layout, cable type, other relevant environmental information)
- Network load at the time of trouble (if known)
- The device history (for example, if you have returned the device before, or if this is a recurring problem)
- Any related RMA (Return Material Authorization) numbers

Subscribe to Product Announcements

You can subscribe to email notifications for product and software release announcements, Field Notices, and Vulnerability Notices.

1. Go to [The Hub](#).
2. In the list of categories, expand the **Product Announcements** list.
3. Select a product for which you would like to receive notifications.
4. Select **Subscribe**.
5. To select additional products, return to the **Product Announcements** list and repeat steps 3 and 4.

You can modify your product selections or unsubscribe at any time.



Command Reference Overview

- [Audience](#) on page 6
- [Structure of this Guide](#) on page 7
- [Product Overview](#) on page 8
- [Software Required](#) on page 8
- [Understanding the Command Syntax](#) on page 9
- [Port Numbering](#) on page 12
- [Line-Editing Keys](#) on page 13
- [Command History](#) on page 14
- [Extreme Networks PoE Devices](#) on page 14

This guide provides details of the command syntax for all Switch Engine commands in this Switch Engine version.

The guide does not provide feature descriptions, explanations of the technologies, or configuration examples. For information about the various features and technologies supported by Extreme Networks switches, see [Switch Engine 32.3 User Guide](#).

This chapter includes the following sections:

- [Audience](#)
- [Structure of this Guide](#)
- [Understanding the Command Syntax](#)
- [Port Numbering](#)
- [Line-Editing Keys](#)
- [Command History](#)

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), Open Shortest Path First (OSPF), and Intermediate System-Intermediate System (IS-IS) concepts.
- Border Gateway Protocol (BGP-4) concepts.
- IP Multicast concepts.
- Protocol Independent Multicast (PIM) concepts.
- Simple Network Management Protocol (SNMP).

Structure of this Guide

This guide documents each Switch Engine command.

Related commands are grouped together and organized into chapters based on their most common usage. The chapters reflect the organization of [Switch Engine 32.3 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.

For each command, the following information is provided:

- **Command Syntax**—The actual syntax of the command. The syntax conventions (the use of braces, for example) are defined in the section [Understanding the Command Syntax](#) on page 9.
- **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 Switch Engine in which the command was introduced, and version(s) where it was modified, if appropriate.
- **Platform Availability**—Platforms on which the command is available.

Product Overview

This table lists the Extreme Networks products that run the Switch Engine software.

Table 1: Switch Engine Switches

Switch Series	Switch Models
ExtremeSwitching 5320	5320-48T-8XE 5320-48P-8XE 5320-24T-8XE 5320-24P-8XE 5320-16P-4XE 5320-16P-4XE-DC
ExtremeSwitching 5420	5420F-8W-16P-4XE 5420F-24P-4XE 5420F-24S-4XE 5420F-24T-4XE 5420F-16MW-32P-4XE 5420F-16W-32P-4XE 5420F-48P-4XE 5420F-48P-4XL 5420F-48T-4XE 5420M-24T-4YE 5420M-24W-4YE 5420M-16MW-32P-4YE 5420M-48T-4YE 5420M-48W-4YE
ExtremeSwitching 5520	5520-24T 5520-24W 5520-48T 5520-48W 5520-12MW-36W 5520-24X 5520-48SE
ExtremeSwitching 5720	5720-24MW 5720-24MXW 5720-48MW 5720-48MXW

Software Required

[ExtremeXOS and Switch Engine Software Support](#) lists the minimum Switch Engine software version required to support each ExtremeSwitching switch model.



Note

The features available on each switch are determined by the installed feature license and optional feature packs. For more information, see the [Feature License Requirements](#) document.

A SummitStack is a combination of up to eight Summit family switches that are connected together.

Understanding the Command Syntax

This section covers the following topics:

- [Access Levels](#)
- [Syntax Symbols](#)
- [Syntax Helper](#) on page 10
- [Object Names](#)
- [Command Shortcuts](#)

Access Levels

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.

Syntax Symbols

You may see a variety of symbols shown as part of the command syntax.

These symbols explain how to enter the command, and you do not type them as part of the command itself. The following table summarizes command syntax symbols.



Note

Switch Engine software does not support the ampersand (&), left angle bracket (<), or right angle bracket (>) because they are reserved characters with special meaning in XML.

Table 2: Command Syntax Symbols

Symbol	Description
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.
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 { time <i>month day year hour min sec</i> } { cancel } { msm-slot-id } { slots <i>slot-number</i> node-address <i>node-address</i> stack-topology { as-standby }} you can specify either a particular date and time combination, or the keyword cancel to cancel a previously scheduled reboot. In this command, 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.
vertical bar	Separates mutually exclusive items in a list, one of which must be entered. For example, in the syntax configure snmp community [readonly readwrite] <i>alphanumeric_string</i> you must specify either the read or write community string in the command. Do not type the vertical bar.

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 also lists any currently configured names that might be used as the next option. In situations where this list might be very long, the syntax helper lists only one line of names, followed by an ellipsis (...) to indicate that there are more names than can be displayed.

Some values (such as the *node-address* used in Summit stack) are lengthy, but limited in number. Switch Engine places these values into a "namespace." This allows command completion on these values.

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 provides 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.

CLI File Name Completion

When entering a command, at the point where a file name could be entered, press **[Tab]** or **?** to display an alphabetically sorted list of possible file names. You can also type part of the file name to display a filtered list of file names matching what you have typed so far.

The following commands support this behavior:

- `cd directory_name`
- `configure access-list aclname [any | ports port_list | vlan vlan_name] {ingress | egress}`
- `configure ip-security dhcp-bindings storage filename name`
- `configure snmp access-profile [access_profile {readonly | readwrite} | [[add rule] [first | [[before | after] previous_rule]]] | delete rule | none]`
- `configure ssh2 access-profile [access_profile | [[add rule] [first | [[before | after] previous_rule]]] | delete rule | none]`
- `configure telnet access-profile [access_profile | [[add rule] [first | [[before | after] previous_rule]]] | delete rule | none]`
- `configure vlan vlan_name udp-profile [profilename | none]`
- `cp old_name new_name`

- create process *name* executable *exe* {start [**auto** | **on-demand**]} {**node** *node*} {**vr** *vr-name*} {**description** *description*} {*arg1* {*arg2* { *arg3* { *arg4* { *arg5* { *arg6* { *arg7* { *arg8* { *arg9* }}}}}}}}}}
- create process *name* **python-module** *python-module* {**start** [**auto** | **on-demand**]} {**node** *node*} {**vr** *vr-name*} {**description** *description*} {*arg1* {*arg2* {*arg3* {*arg4* {*arg5* {*arg6* {*arg7* {*arg8* {*arg9*}}}}}}}}}
- edit policy *filename*
- enable license file *filename*
- enable ssh2 {**access-profile** [*access_profile* | **none**]} {**port** *tcp_port_number*} {**vr** [*vr_name* | **all** | **default**]}
- load script *filename* {**arg1**} {**arg2**} ... {**arg9**}
- ls *file_name*
- mkdir *directory_name*
- mv *old_name* *new_name*
- scp2 {**cipher** *cipher*} {**mac** *mac*} {**compression** [**on** | **off**]} {**port** *portnum*} {**vr** *vr_name*} user [*hostname* | *ipaddress*]:*remote_file* *local_file*
- show ssl {[**trusted-ca** | **ocsp-signature-ca**] [*file_name* | **all**]} {**manufacturing**} {**certificate** | **detail**}
- tftp [**ip-address** | **host-name**] { **-v** *vr_name* } { **-b** *block_size* } [**-g** | **-p**] [**-l** *local-file* { **-r** *remote-file* } | **-r** *remote-file* { **-l** *local-file* }]
- tftp get [**ip-address** | **host-name**] { **vr** *vr_name* } { **block-size** *block_size* } *remote-file* *local-file* {**force-overwrite**}
- tftp put [*ip-address* | *host-name*] {**vr** *vr_name*} {**block-size** *block_size*}*local-file* { *remote-file*}
- rm *file_name*
- rmdir *directory_name*
- run script *filename* {**arg1**} {**arg2**} ... {**arg9**}
- unconfigure ssl certificate [**trusted-ca** | **ocsp-signature-ca**] [*file_name* | **all**]

Object Names

All named components within a category of the switch configuration, such as VLAN (Virtual LAN), must be given a unique object name.

Object names must begin with an alphabetical character and may contain alphanumeric characters and underscores (_), but they cannot contain spaces. The maximum allowed length for a name is 32 characters.

Object names can be reused across categories (for example, *STPD (Spanning Tree Domain)* and VLAN names). If the software encounters any ambiguity in the components within your command, it generates a message requesting that you clarify the object you specified.

**Note**

If you use the same name across categories, Extreme Networks recommends that you specify the identifying keyword as well as the actual name. If you do not use the keyword, the system may return an error message.

Reserved Keywords

Keywords such as `vlan`, *STP (Spanning Tree Protocol)*, and other 2nd level keywords, are determined to be reserved keywords and cannot be used as object names. This restriction applies to the specific word (`vlan`) only, while expanded versions (`vlan2`) can be used.

A complete list of the reserved keywords for Switch Engine is found in the “Reserved Keywords” section of the *Switch Engine 32.3 User Guide*. Any keyword that is not on this list can be used as an object name.

Command Shortcuts

Components are typically named using the `create` command.

When you enter a command to configure a named component, you do not need to use the keyword of the component. For example, to create a *VLAN*, enter a 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 (unless you used the same name for another category such as *STPD* or *EAPS (Extreme Automatic Protection Switching)*).

For example, instead of entering the command:

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

you could enter the following shortcut:

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

Port Numbering

Commands that require you to enter one or more port numbers use the parameter *port_list* in the syntax.

The available variables differ on a stand-alone switch and SummitStack.

**Note**

The keyword `all` acts on all possible ports; it continues on all ports even if one port in the sequence fails.

Stand-alone Switch Numerical Ranges

On ExtremeSwitching switches, the port number is simply noted by the physical port number.

Separate the port numbers by a dash to enter a range of contiguous numbers, and separate the numbers by a comma to enter a range of non-contiguous numbers:

- x-y—Specifies a contiguous series of ports on a stand-alone switch.
- x,y—Specifies a non-contiguous series of ports on a stand-alone switch.
- x-y,a,d—Specifies a contiguous series of ports and a non-contiguous series of ports on a stand-alone switch.
- Port:Channel—For ExtremeSwitching 5720 channelized ports. For example, 49:4 maps to Port 49, Channel 4.

SummitStack Numerical Ranges

On SummitStack switches, the port number is a combination of the slot number and the port number.

The nomenclature for the port number is as follows: `slot:port`

For example, if there is a switch in slot 2 of the stack with a total of four ports, the following ports are valid:

- 2:1
- 2:2
- 2:3
- 2:4

You can also use wildcard combinations (*) to specify port combinations.

The following wildcard combinations are allowed:

- slot:*—Specifies all ports on a particular switch in the stack.
- slot:x-slot:y—Specifies a contiguous series of ports on a range of switches in the stack.
- slot:x-y—Specifies a contiguous series of ports on a particular switch in the stack.
- slota:x-slotb:y—Specifies a contiguous series of ports on a SummitStack node and end on another node.
- Slot:Port:Channel—For ExtremeSwitching 5720 channelized ports. For example, 2:49:4 maps to Slot 2, Port 49, Channel 4.

Line-Editing Keys

Table 3 describes the line-editing keys available using the CLI.

Table 3: Line-Editing Keys

Key(s)	Description
Left arrow or [Ctrl] + B	Moves the cursor one character to the left.
Right arrow or [Ctrl] + F	Moves the cursor one character to the right.
[Ctrl] + H or Backspace	Deletes character to left of cursor and shifts remainder of line to left.

Table 3: Line-Editing Keys (continued)

Key(s)	Description
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.
[Ctrl] + A	Moves cursor to first character in line.
[Ctrl] + E	Moves cursor to last character in line.
[Ctrl] + L	Clears screen and moves cursor to beginning of line.
[Ctrl] + P or Up Arrow	Displays previous command in command history buffer and places cursor at end of command.
[Ctrl] + N or Down Arrow	Displays next command in command history buffer and places cursor at end of command.
[Ctrl] + U	Clears all characters typed from cursor to beginning of line.
[Ctrl] + W	Deletes previous word.
[Ctrl] + C	Interrupts the current CLI command execution.

Command History

The operating system "remembers" all the commands you enter.

You can display a list of these commands by using the following command:

```
history
```

If you use a command more than once, consecutively, the history will list only the first instance.

Extreme Networks PoE Devices

Following is a list of the Extreme Networks devices that support PoE+ and PoE++ and the minimum required software:

PoE+

- ExtremeSwitching 5420F-8W-16P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420F-24P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420F-16MW-32P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420F-16W-32P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420F-48P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420F-48P-4XL—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420M-16MW-32P-4YE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5320-48P-8XE—ExtremeXOS 31.6 and later.
- ExtremeSwitching 5320-24P-8XE—ExtremeXOS 31.6 and later.

- ExtremeSwitching 5320-16P-4XE—ExtremeXOS 31.7 and later.
- ExtremeSwitching 5320-16P-4XE-DC—ExtremeXOS 31.7 and later.

PoE++

- ExtremeSwitching 5520-24W—ExtremeXOS 31.1 and later.
- ExtremeSwitching 5520-48W—ExtremeXOS 31.1 and later.
- ExtremeSwitching 5520-12MW-36W—ExtremeXOS 31.1 and later.
- ExtremeSwitching 5420F-8W-16P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420F-16MW-32P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420F-16W-32P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420M-24W-4YE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420M-16MW-32P-4YE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420M-48W-4YE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5720-24MW—Switch Engine 32.1 and later.
- ExtremeSwitching 5720-24MXW—Switch Engine 32.1 and later.
- ExtremeSwitching 5720-48MW—Switch Engine 32.1 and later.
- ExtremeSwitching 5720-48MXW—Switch Engine 32.1 and later.



Commands

[alias](#) on page 87
[cat](#) on page 89
[cd](#) on page 91
[check policy attribute](#) on page 91
[check policy](#) on page 93
[clear access-list counter](#) on page 94
[clear access-list meter](#) on page 95
[clear account lockout](#) on page 96
[clear bgp flap-statistics](#) on page 97
[clear bgp neighbor counters](#) on page 99
[clear bootprelay ipv6 prefix-delegation snooping](#) on page 100
[clear cdp counters](#) on page 101
[clear cdp neighbor](#) on page 101
[clear counters](#) on page 102
[clear counters bfd](#) on page 103
[clear counters bfd missed-hellos](#) on page 104
[clear counters cfm segment all](#) on page 105
[clear counters cfm segment all frame-delay](#) on page 108
[clear counters cfm segment all frame-loss](#) on page 110
[clear counters cfm segment frame-delay](#) on page 113
[clear counters cfm segment frame-loss mep](#) on page 114
[clear counters cfm segment frame-loss](#) on page 116
[clear counters cfm segment](#) on page 117
[clear counters cfm session missed-hellos](#) on page 119
[clear counters edp](#) on page 120
[clear counters erps](#) on page 121
[clear counters mpls](#) on page 121
[clear counters fdb mac-tracking](#) on page 122
[clear counters flowmon](#) on page 123
[clear counters identity-management](#) on page 124
[clear counters iparp](#) on page 125
[clear counters l2vpn](#) on page 126
[clear counters mpls ldp](#) on page 127
[clear counters mpls rsvp-te](#) on page 127

[clear counters mpls static lsp](#) on page 128
[clear counters policy](#) on page 129
[clear counters ports](#) on page 130
[clear counters ports protocol filter](#) on page 131
[clear counters stp](#) on page 132
[clear counters virtual-network](#) on page 133
[clear counters virtual-network remote-endpoint](#) on page 134
[clear counters vpls](#) on page 135
[clear counters vr](#) on page 135
[clear counters vrrp](#) on page 136
[clear counters wred ecn](#) on page 137
[clear counters xml-notification](#) on page 138
[clear cpu-monitoring](#) on page 139
[clear dns cache](#) on page 139
[clear dns cache analytics entries](#) on page 140
[clear eaps counters](#) on page 141
[clear elrp counters](#) on page 142
[clear elsm ports auto-restart](#) on page 143
[clear elsm ports counters](#) on page 144
[clear esrp counters](#) on page 145
[clear esrp neighbor](#) on page 146
[clear esrp sticky](#) on page 147
[clear ethernet oam counters](#) on page 148
[clear fdb](#) on page 149
[clear fdb vpls](#) on page 150
[clear igmp counters](#) on page 151
[clear igmp group](#) on page 152
[clear igmp snooping](#) on page 153
[clear inline-power stats ports](#) on page 154
[clear ip nat counters vlan](#) on page 155
[clear iparp](#) on page 156
[clear ip-security anomaly-protection notify cache](#) on page 157
[clear ip-security arp validation violations](#) on page 158
[clear ip-security dhcp-snooping entries](#) on page 158
[clear ip-security source-ip-lockdown entries ports](#) on page 159
[clear ipv6 dad](#) on page 160
[clear isis counters](#) on page 161
[clear isis counters area](#) on page 161
[clear isis counters vlan](#) on page 162
[clear l2pt counters rtep](#) on page 163
[clear l2pt counters vlan](#) on page 164
[clear l2pt counters vman](#) on page 165

[clear l2pt counters vpls](#) on page 166
[clear lacp counters](#) on page 167
[clear lldp neighbors](#) on page 168
[clear log](#) on page 168
[clear log counters](#) on page 169
[clear mac-locking station](#) on page 171
[clear macsec counters](#) on page 171
[clear meter out-of-profile](#) on page 172
[clear mld counters](#) on page 174
[clear mld group](#) on page 174
[clear mld snooping](#) on page 175
[clear msdp counters](#) on page 176
[clear msdp sa-cache](#) on page 178
[clear neighbor-discovery cache](#) on page 179
[clear netlogin state](#) on page 180
[clear netlogin state agent](#) on page 181
[clear netlogin state mac-address](#) on page 181
[clear network-clock gtp counters](#) on page 182
[clear nodealias](#) on page 183
[clear ospf counters](#) on page 184
[clear ospfv3 counters](#) on page 185
[clear pim cache](#) on page 186
[clear pim snooping](#) on page 187
[clear port rate-limit flood](#) on page 188
[clear ports link-flap-detection counters](#) on page 189
[clear ports link-flap-detection status](#) on page 190
[clear port rate-limit flood](#) on page 191
[clear process group statistics](#) on page 192
[clear rip counters](#) on page 193
[clear ripng counters](#) on page 194
[clear screen](#) on page 194
[clear session](#) on page 195
[clear slot](#) on page 196
[clear snmp notification-log](#) on page 197
[clear stpd ports](#) on page 198
[clear switch bluetooth](#) on page 199
[clear vm storage](#) on page 200
[clear vlan dhcp-address-allocation](#) on page 201
[configure access-list](#) on page 202
[configure access-list action-resolution highest-priority](#) on page 203
[configure access-list action-resolution multiple](#) on page 204
[configure access-list add](#) on page 205

[configure access-list delete](#) on page 207
[configure access-list network-zone](#) on page 208
[configure access-list rule-compression port-counters](#) on page 209
[configure access-list vlan-acl-precedence](#) on page 210
[configure access-list width](#) on page 211
[configure access-list zone](#) on page 212
[configure account](#) on page 213
[configure account encrypted](#) on page 215
[configure account password-policy char-validation](#) on page 216
[configure account password-policy history](#) on page 217
[configure account password-policy lockout-on-login-failures](#) on page 218
[configure account password-policy lockout-time-period](#) on page 220
[configure account password-policy max-age](#) on page 220
[configure account password-policy min-age](#) on page 222
[configure account password-policy min-different-characters](#) on page 223
[configure account password-policy min-length](#) on page 224
[configure account privilege](#) on page 225
[configure auto-peering oneconfig bootprelay](#) on page 226
[configure auto-peering oneconfig id](#) on page 227
[configure auto-peering one-config iproute](#) on page 228
[configure auto-peering one-config nsi-id](#) on page 230
[configure auto-peering oneconfig overlay](#) on page 231
[configure auto-peering one-config password](#) on page 232
[configure auto-peering one-config remote id](#) on page 233
[configure automation edge connect/disconnect](#) on page 234
[configure banner](#) on page 235
[configure bfd vlan](#) on page 237
[configure bfd vlan authentication](#) on page 238
[configure bgp add aggregate-address](#) on page 239
[configure bgp add confederation-peer sub-AS-number](#) on page 241
[configure bgp add network](#) on page 242
[configure bgp as-display-format](#) on page 243
[configure bgp as-number](#) on page 244
[configure bgp cluster-id](#) on page 245
[configure bgp confederation-id](#) on page 246
[configure bgp delete aggregate-address](#) on page 247
[configure bgp delete confederation-peer sub-AS-number](#) on page 249
[configure bgp delete network](#) on page 250
[configure bgp evpn ignore-as](#) on page 251
[configure bgp evpn instance rd](#) on page 252
[configure bgp evpn instance route-target](#) on page 253
[configure bgp evpn instance vxlan](#) on page 254

[configure bgp evpn l3vni](#) on page 255
[configure bgp export shutdown-priority](#) on page 256
[configure bgp import-policy](#) on page 257
[configure bgp local-preference](#) on page 258
[configure bgp maximum-as-path-length](#) on page 259
[configure bgp maximum-paths](#) on page 260
[configure bgp med](#) on page 261
[configure bgp neighbor allowas-in](#) on page 262
[configure bgp neighbor alternate-local-as](#) on page 264
[configure bgp neighbor bfd](#) on page 265
[configure bgp neighbor connect-retry](#) on page 266
[configure bgp neighbor dampening](#) on page 267
[configure bgp neighbor description](#) on page 269
[configure bgp neighbor dont-allowas-in](#) on page 270
[configure bgp neighbor maximum-prefix](#) on page 272
[configure bgp neighbor next-hop-self](#) on page 274
[configure bgp neighbor no-dampening](#) on page 276
[configure bgp neighbor password](#) on page 277
[configure bgp neighbor peer-group](#) on page 279
[configure bgp neighbor route-policy](#) on page 280
[configure bgp neighbor route-reflector-client](#) on page 282
[configure bgp neighbor send-community](#) on page 284
[configure bgp neighbor shutdown-priority](#) on page 285
[configure bgp neighbor soft-reset](#) on page 286
[configure bgp neighbor source-interface](#) on page 288
[configure bgp neighbor timer](#) on page 289
[configure bgp neighbor weight](#) on page 290
[configure bgp peer-group allowas-in](#) on page 292
[configure bgp peer-group connect-retry](#) on page 293
[configure bgp peer-group dampening](#) on page 294
[configure bgp peer-group dont-allowas-in](#) on page 296
[configure bgp peer-group maximum-prefix](#) on page 298
[configure bgp peer-group next-hop-self](#) on page 300
[configure bgp peer-group no-dampening](#) on page 301
[configure bgp peer-group password](#) on page 303
[configure bgp peer-group remote-AS-number](#) on page 304
[configure bgp peer-group route-policy](#) on page 305
[configure bgp peer-group route-reflector-client](#) on page 306
[configure bgp peer-group send-community](#) on page 307
[configure bgp peer-group soft-reset](#) on page 309
[configure bgp peer-group source-interface](#) on page 311
[configure bgp peer-group timer](#) on page 312

[configure bgp peer-group weight](#) on page 313

[configure bgp restart address-family](#) on page 314

[configure bgp restart restart-time](#) on page 315

[configure bgp restart stale-route-time](#) on page 316

[configure bgp restart update-delay](#) on page 317

[configure bgp restart](#) on page 318

[configure bgp routerid](#) on page 319

[configure bgp soft-reconfiguration](#) on page 320

[configure bootprelay](#) on page 321

[configure bootprelay add](#) on page 322

[configure bootprelay delete](#) on page 323

[configure bootprelay dhcp-agent information check](#) on page 324

[configure bootprelay dhcp-agent information circuit-id port-information](#) on page 325

[configure bootprelay dhcp-agent information circuit-id vlan-information](#) on page 326

[configure bootprelay dhcp-agent information option](#) on page 327

[configure bootprelay dhcp-agent information policy](#) on page 328

[configure bootprelay dhcp-agent information remote-id](#) on page 329

[configure bootprelay dhcp-agent source-vlan](#) on page 329

[configure bootprelay include-secondary](#) on page 331

[configure bootprelay ipv6 option interface-id](#) on page 332

[configure bootprelay ipv6 option remote-id](#) on page 333

[configure bootprelay ipv6 prefix-delegation snooping add](#) on page 334

[configure bootprelay ipv6 prefix-delegation snooping](#) on page 335

[configure bootprelay vlan include-secondary](#) on page 336

[configure cdp cos-extend ports](#) on page 337

[configure cdp device-id](#) on page 338

[configure cdp frequency](#) on page 339

[configure cdp hold-time](#) on page 340

[configure cdp management-address](#) on page 340

[configure cdp power-available ports](#) on page 342

[configure cdp trust-extend ports](#) on page 343

[configure cdp voip-vlan ports](#) on page 343

[configure cfm domain add association integer](#) on page 345

[configure cfm domain add association meg](#) on page 346

[configure cfm domain add association string](#) on page 347

[configure cfm domain add association vlan-id](#) on page 348

[configure cfm domain add association vpn-id oui index](#) on page 349

[configure cfm domain association add remote-mep](#) on page 350

[configure cfm domain association add](#) on page 350

[configure cfm domain association delete remote-mep](#) on page 352

[configure cfm domain association delete](#) on page 353

[configure cfm domain association destination-mac-type](#) on page 354

[configure cfm domain association end-point add group](#) on page 355

[configure cfm domain association end-point delete group](#) on page 356

[configure cfm domain association end-point transmit-interval](#) on page 357

[configure cfm domain association ports end-point ccm](#) on page 358

[configure cfm domain association ports end-point mepid](#) on page 359

[configure cfm domain association ports end-point sender-id-ipaddress](#) on page 360

[configure cfm domain association ports end-point](#) on page 361

[configure cfm domain association remote-mep mac-address](#) on page 362

[configure cfm domain delete association](#) on page 363

[configure cfm domain md-level](#) on page 363

[configure cfm group add rmep](#) on page 364

[configure cfm group delete rmep](#) on page 365

[configure cfm segment add domain association](#) on page 366

[configure cfm segment delete domain association](#) on page 367

[configure cfm segment dot1p](#) on page 367

[configure cfm segment frame-delay dot1p](#) on page 368

[configure cfm segment frame-delay window](#) on page 369

[configure cfm segment frame-delay/frame-loss transmit interval](#) on page 370

[configure cfm segment frame-loss consecutive](#) on page 371

[configure cfm segment frame-loss dot1p](#) on page 371

[configure cfm segment frame-loss mep](#) on page 372

[configure cfm segment frame-loss ses-threshold](#) on page 373

[configure cfm segment frame-loss window](#) on page 374

[configure cfm segment threshold](#) on page 375

[configure cfm segment timeout](#) on page 376

[configure cfm segment transmit-interval](#) on page 376

[configure cfm segment window](#) on page 377

[configure cli](#) on page 378

[configure cli journal](#) on page 380

[configure cli max-failed-logins](#) on page 381

[configure cli max-sessions](#) on page 381

[configure cli mode](#) on page 382

[configure cli mode scripting](#) on page 383

[configure cli moved-keywords](#) on page 384

[configure cli password prompting-only](#) on page 385

[configure cli script path](#) on page 386

[configure cli script timeout](#) on page 387

[configure cos-index](#) on page 388

[configure database add server](#) on page 389

[configure database delete server](#) on page 390

[configure database max-retry-interval](#) on page 391

[configure database server password](#) on page 392

[configure debug core-dumps](#) on page 393
[configure dhcp ipv6 client identifier-type](#) on page 394
[configure diagnostics privilege](#) on page 395
[configure diffserv examination code-point qosprofile](#) on page 396
[configure diffserv replacement code-point](#) on page 397
[configure dns cache analytics \[add | delete\] protected-client](#) on page 399
[configure dns cache add | delete name-server](#) on page 400
[configure dns cache analytics](#) on page 401
[configure dns-client add](#) on page 402
[configure dns-client default-domain](#) on page 403
[configure dns-client delete](#) on page 404
[configure dos-protect acl-expire](#) on page 405
[configure dos-protect interval](#) on page 406
[configure dos-protect trusted ports](#) on page 407
[configure dos-protect type l3-protect alert-threshold](#) on page 408
[configure dos-protect type l3-protect notify-threshold](#) on page 408
[configure dot1p type](#) on page 409
[configure eaps add control vlan](#) on page 411
[configure eaps add protected vlan](#) on page 412
[configure eaps cfm](#) on page 413
[configure eaps config-warnings off](#) on page 414
[configure eaps config-warnings on](#) on page 415
[configure eaps delete control vlan](#) on page 416
[configure eaps delete protected vlan](#) on page 417
[configure eaps failtime expiry-action](#) on page 418
[configure eaps failtime](#) on page 419
[configure eaps fast-convergence](#) on page 420
[configure eaps hello-pdu-egress](#) on page 421
[configure eaps hellotime](#) on page 422
[configure eaps mode](#) on page 423
[configure eaps multicast add-ring-ports](#) on page 425
[configure eaps multicast send-igmp-query](#) on page 426
[configure eaps multicast temporary-flooding duration](#) on page 427
[configure eaps multicast temporary-flooding](#) on page 428
[configure eaps name](#) on page 429
[configure eaps port](#) on page 430
[configure eaps priority](#) on page 431
[configure eaps shared-port common-path-timers](#) on page 432
[configure eaps shared-port link-id](#) on page 433
[configure eaps shared-port mode](#) on page 434
[configure eaps shared-port segment-timers expiry-action](#) on page 435
[configure eaps shared-port segment-timers health-interval](#) on page 437

[configure eaps shared-port segment-timers timeout](#) on page 437

[configure edp advertisement-interval](#) on page 438

[configure elrp-client dynamic-vlans](#) on page 439

[configure elrp-client dynamic-vlans action](#) on page 440

[configure elrp-client dynamic-vlans client/uplink ports/remote-endpoints vxlan](#) on page 442

[configure elrp-client disable ports](#) on page 443

[configure elrp-client hardware-assist](#) on page 445

[configure elrp-client inter-vlan-loop-detection](#) on page 446

[configure elrp-client one-shot](#) on page 447

[configure elrp-client periodic](#) on page 448

[configure elsm ports hellotime](#) on page 451

[configure elsm ports hold-threshold](#) on page 452

[configure elsm ports uptimer-threshold](#) on page 453

[configure erps add control vlan](#) on page 454

[configure erps add protected vlan](#) on page 455

[configure erps control-mac](#) on page 456

[configure erps cfm port group](#) on page 457

[configure erps cfm protection group](#) on page 458

[configure erps delete control vlan](#) on page 459

[configure erps delete protected vlan](#) on page 460

[configure erps dynamic-state](#) on page 461

[configure erps name](#) on page 462

[configure erps neighbor port](#) on page 462

[configure erps notify-topology-change](#) on page 463

[configure erps protection-port](#) on page 464

[configure erps revert](#) on page 465

[configure erps ring-ports east | west](#) on page 466

[configure erps subring-mode](#) on page 466

[configure erps sub-ring](#) on page 467

[configure erps timer guard](#) on page 468

[configure erps timer hold-off](#) on page 469

[configure erps timer periodic](#) on page 470

[configure erps timer wait-to-block](#) on page 470

[configure erps timer wait-to-restore](#) on page 471

[configure erps topology-change](#) on page 472

[configure esrp add elrp-poll ports](#) on page 473

[configure esrp add master](#) on page 474

[configure esrp add member](#) on page 475

[configure esrp add track-environment](#) on page 476

[configure esrp add track-iproute](#) on page 477

[configure esrp add track-ping](#) on page 478

[configure esrp add track-vlan](#) on page 479
[configure esrp aware add selective-forward-ports](#) on page 480
[configure esrp aware delete selective-forward-ports](#) on page 481
[configure esrp delete elrp-poll ports](#) on page 482
[configure esrp delete master](#) on page 483
[configure esrp delete member](#) on page 484
[configure esrp delete track-environment](#) on page 484
[configure esrp delete track-iproute](#) on page 485
[configure esrp delete track-ping](#) on page 486
[configure esrp delete track-vlan](#) on page 487
[configure esrp domain-id](#) on page 487
[configure esrp election-policy](#) on page 488
[configure esrp elrp-master-poll disable](#) on page 492
[configure esrp elrp-master-poll enable](#) on page 493
[configure esrp elrp-premaster-poll disable](#) on page 494
[configure esrp elrp-premaster-poll enable](#) on page 495
[configure esrp group](#) on page 496
[configure esrp mode](#) on page 497
[configure esrp name](#) on page 498
[configure esrp ports mode](#) on page 499
[configure esrp ports no-restart](#) on page 500
[configure esrp ports restart](#) on page 501
[configure esrp ports weight](#) on page 501
[configure esrp priority](#) on page 502
[configure esrp timer hello](#) on page 503
[configure esrp timer neighbor](#) on page 504
[configure esrp timer neutral](#) on page 506
[configure esrp timer premaster](#) on page 507
[configure esrp timer restart](#) on page 508
[configure failsafe-account](#) on page 509
[configure fabric attach management-vlan](#) on page 510
[configure fabric attach management-vlan ports](#) on page 512
[configure fabric attach ports](#) on page 513
[configure fabric attach port authentication](#) on page 513
[configure fabric attach uplink](#) on page 515
[configure fabric attach zero-touch-client](#) on page 516
[configure fdb agingtime](#) on page 517
[configure fdb mac-tracking ports](#) on page 518
[configure fdb static-mac-move packets](#) on page 519
[configure fdb vlan vxlan](#) on page 520
[configure flow-redirect add nexthop](#) on page 521
[configure flow-redirect delete nexthop](#) on page 522

[configure flow-redirect health-check](#) on page 523
[configure flow-redirect nexthop](#) on page 524
[configure flow-redirect no-active](#) on page 525
[configure flow-redirect vr](#) on page 526
[configure flowmon collector](#) on page 527
[configure flowmon group](#) on page 528
[configure flowmon group collector](#) on page 529
[configure flowmon group key](#) on page 530
[configure flowmon key ipv4](#) on page 531
[configure flowmon key ipv6](#) on page 532
[configure forwarding internal-tables](#) on page 534
[configure forwarding flow-control fabric](#) on page 536
[configure forwarding hash-algorithm](#) on page 537
[configure forwarding hash-recursion-level](#) on page 538
[configure forwarding ipmc all](#) on page 539
[configure forwarding ipmc compression](#) on page 540
[configure forwarding ipmc llmnr](#) on page 541
[configure forwarding ipmc local-network-range](#) on page 542
[configure forwarding ipmc lookup-key](#) on page 543
[configure forwarding ipmc mdns](#) on page 544
[configure forwarding ipmc upnp](#) on page 545
[configure forwarding L2-protocol fast-convergence](#) on page 546
[configure forwarding rate-limit overhead-bytes](#) on page 547
[configure forwarding sharing](#) on page 548
[configure forwarding suppression filters](#) on page 549
[configure forwarding vpex ipmc replication](#) on page 550
[configure forwarding vpex vlan-port-filter](#) on page 551
[configure identity-management role](#) on page 552
[configure identity-management role-based-vlan](#) on page 553
[configure identity-management access-list](#) on page 554
[configure identity-management blacklist](#) on page 555
[configure identity-management database memory-size](#) on page 557
[configure identity-management detection](#) on page 558
[configure identity-management greylist](#) on page 560
[configure identity-management kerberos snooping aging time](#) on page 561
[configure identity-management kerberos snooping force-aging time](#) on page 562
[configure identity-management kerberos snooping forwarding](#) on page 563
[configure identity-management kerberos snooping server](#) on page 564
[configure identity-management list-precedence](#) on page 565
[configure identity-management ports](#) on page 567
[configure identity-management role add child-role](#) on page 568
[configure identity-management role add dynamic-rule](#) on page 569

[configure identity-management role add policy](#) on page 570

[configure identity-management role delete child-role](#) on page 571

[configure identity-management role delete dynamic-rule](#) on page 571

[configure identity-management role delete policy](#) on page 572

[configure identity-management role match-criteria inheritance](#) on page 573

[configure identity-management role priority](#) on page 574

[configure identity-management stale-entry aging-time](#) on page 575

[configure identity-management whitelist](#) on page 578

[configure cli idle-timeout](#) on page 580

[configure igmp](#) on page 580

[configure igmp router-alert receive-required](#) on page 582

[configure igmp router-alert transmit](#) on page 583

[configure igmp snooping filters](#) on page 584

[configure igmp snooping flood-list](#) on page 585

[configure igmp snooping leave-timeout](#) on page 587

[configure igmp snooping timer](#) on page 588

[configure igmp snooping vlan ports add dynamic group](#) on page 589

[configure igmp snooping vlan ports add static group](#) on page 590

[configure igmp snooping vlan ports add static router](#) on page 592

[configure igmp snooping vlan ports delete static group](#) on page 593

[configure igmp snooping vlan ports delete static router](#) on page 594

[configure igmp snooping vlan ports filter](#) on page 595

[configure igmp snooping vlan ports set join-limit](#) on page 596

[configure igmp ssm-map add](#) on page 597

[configure igmp ssm-map delete](#) on page 598

[configure inline-power classification](#) on page 599

[configure inline-power detection ports](#) on page 601

[configure inline-power disconnect-precedence](#) on page 603

[configure inline-power label ports](#) on page 605

[configure inline-power operator-limit ports](#) on page 606

[configure inline-power priority ports](#) on page 608

[configure inline-power usage-threshold](#) on page 610

[configure ip anycast mac](#) on page 612

[configure ip nat add vlan](#) on page 613

[configure ip nat aging-time](#) on page 614

[configure ip nat delete vlan](#) on page 615

[configure ip nat rule destination](#) on page 616

[configure ip nat rule destination protocol](#) on page 617

[configure ip nat rule egress](#) on page 618

[configure ip nat rule monitor](#) on page 619

[configure ip nat rule name](#) on page 620

[configure ip nat rule source](#) on page 621

[configure iparp add proxy](#) on page 622
[configure iparp add](#) on page 624
[configure iparp delete proxy](#) on page 624
[configure iparp delete](#) on page 625
[configure ip-arp fast-convergence](#) on page 626
[configure iparp locktime](#) on page 627
[configure iparp max_entries](#) on page 628
[configure iparp max_pending_entries](#) on page 629
[configure iparp max_proxy_entries](#) on page 630
[configure iparp proxy reachable | entry-required](#) on page 631
[configure iparp reachable-time](#) on page 632
[configure iparp retransmit-time](#) on page 633
[configure iparp timeout](#) on page 634
[configure ipforwarding originated-packets](#) on page 635
[configure ipmcforwarding](#) on page 636
[configure ipmroute add](#) on page 637
[configure ipmroute delete](#) on page 638
[configure ip-mtu vlan](#) on page 639
[configure iproute add blackhole ipv4 default](#) on page 640
[configure iproute add blackhole ipv6 default](#) on page 641
[configure iproute add blackhole](#) on page 642
[configure iproute add default](#) on page 643
[configure iproute add \(IPv4\)](#) on page 644
[configure iproute add \(IPv6\)](#) on page 647
[configure iproute add lsp](#) on page 648
[configure iproute add \(Multicast\)](#) on page 649
[configure iproute add protection](#) on page 651
[configure iproute delete](#) on page 652
[configure iproute delete blackhole](#) on page 653
[configure iproute delete blackhole ipv4 default](#) on page 654
[configure iproute delete blackhole ipv6 default](#) on page 655
[configure iproute delete default](#) on page 656
[configure iproute ipv6 priority](#) on page 657
[configure iproute priority](#) on page 659
[configure iproute reserved-entries](#) on page 662
[configure iproute protection ping interval](#) on page 664
[configure iproute sharing hash-algorithm crc](#) on page 665
[configure iproute sharing max-gateways](#) on page 666
[configure ip-security anomaly-protection icmp ipv4-max-size](#) on page 668
[configure ip-security anomaly-protection icmp ipv6-max-size](#) on page 668
[configure ip-security anomaly-protection notify cache](#) on page 669
[configure ip-security anomaly-protection notify rate limit](#) on page 670

[configure ip-security anomaly-protection notify rate window](#) on page 671

[configure ip-security anomaly-protection notify trigger off](#) on page 671

[configure ip-security anomaly-protection notify trigger on](#) on page 672

[configure ip-security anomaly-protection tcp](#) on page 673

[configure ip-security dhcp-bindings add](#) on page 674

[configure ip-security dhcp-bindings delete](#) on page 675

[configure ip-security dhcp-bindings storage filename](#) on page 676

[configure ip-security dhcp-bindings storage location](#) on page 677

[configure ip-security dhcp-bindings storage](#) on page 678

[configure ip-security dhcp-snooping information check](#) on page 679

[configure ip-security dhcp-snooping information circuit-id port-information port](#) on page 680

[configure ip-security dhcp-snooping information circuit-id vlan-information](#) on page 680

[configure ip-security dhcp-snooping information option](#) on page 681

[configure ip-security dhcp-snooping information policy](#) on page 682

[configure ip-security dhcp-snooping information remote-id](#) on page 683

[configure ipv6 dad](#) on page 684

[configure ipv6 hop-limit](#) on page 685

[configure iqagent http-proxy](#) on page 686

[configure iqagent server](#) on page 687

[configure irdp](#) on page 688

[configure isis add vlan](#) on page 689

[configure isis area add area-address](#) on page 690

[configure isis area add summary-address](#) on page 691

[configure isis area area-password](#) on page 693

[configure isis area delete area-address](#) on page 694

[configure isis area delete summary-address](#) on page 694

[configure isis area domain-password](#) on page 695

[configure isis area interlevel-filter level 1-to-2](#) on page 696

[configure isis area interlevel-filter level 2-to-1](#) on page 697

[configure isis area is-type level](#) on page 698

[configure isis area metric-style](#) on page 699

[configure isis area overload-bit on-startup](#) on page 701

[configure isis area system-id](#) on page 702

[configure isis area timer lsp-gen-interval](#) on page 703

[configure isis area timer lsp-refresh-interval](#) on page 704

[configure isis area timer max-lsp-lifetime](#) on page 705

[configure isis area timer restart](#) on page 706

[configure isis area timer spf-interval](#) on page 707

[configure isis area topology-mode](#) on page 707

[configure isis circuit-type](#) on page 709

[configure isis delete vlan](#) on page 710

[configure isis hello-multiplier](#) on page 711
[configure isis import-policy](#) on page 712
[configure isis link-type](#) on page 713
[configure isis mesh](#) on page 714
[configure isis metric](#) on page 715
[configure isis password vlan](#) on page 716
[configure isis priority](#) on page 717
[configure isis restart grace-period](#) on page 718
[configure isis restart](#) on page 718
[configure isis timer csnp-interval](#) on page 720
[configure isis timer hello-interval](#) on page 720
[configure isis timer lsp-interval](#) on page 721
[configure isis timer restart-hello-interval](#) on page 722
[configure isis timer retransmit-interval](#) on page 723
[configure isis wide-metric](#) on page 724
[configure jumbo-frame-size](#) on page 725
[configure keychain accept-tolerance](#) on page 726
[configure keychain add key](#) on page 727
[configure keychain delete key](#) on page 728
[configure keychain key active-lifetime](#) on page 729
[configure keychain key hash-algorithm](#) on page 730
[configure l2pt profile add profile](#) on page 731
[configure l2pt profile delete profile](#) on page 733
[configure l2vpn add peer](#) on page 734
[configure l2vpn add service](#) on page 735
[configure l2vpn delete peer](#) on page 737
[configure l2vpn delete service](#) on page 738
[configure l2vpn health-check vccv](#) on page 739
[configure l2vpn peer mpls lsp](#) on page 740
[configure l2vpn peer](#) on page 741
[configure l2vpn vpls add peer ipaddress](#) on page 743
[configure vpls add service](#) on page 744
[configure l2vpn vpls peer static-pw](#) on page 745
[configure l2vpn vpls redundancy](#) on page 746
[configure l2vpn vpws add peer ipaddress](#) on page 748
[configure l2vpn vpws peer static-pw](#) on page 750
[configure l2vpn](#) on page 751
[configure lacp member-port priority](#) on page 752
[configure ldap domain](#) on page 754
[configure ldap domain add server](#) on page 754
[configure ldap domain base-dn](#) on page 756
[configure ldap domain bind-user](#) on page 757

[configure ldap domain delete server](#) on page 758
[configure ldap domain netlogin](#) on page 759
[configure ldap hierarchical-search-oid](#) on page 761
[configure lldp management-address](#) on page 761
[configure lldp med fast-start repeat-count](#) on page 763
[configure lldp ports dcbx add application](#) on page 764
[configure lldp ports dcbx delete application](#) on page 765
[configure lldp ports management-address](#) on page 767
[configure lldp ports port-description](#) on page 767
[configure lldp ports system-capabilities](#) on page 768
[configure lldp ports system-description](#) on page 769
[configure lldp ports system-name](#) on page 770
[configure lldp ports vendor-specific avaya-extreme call-server](#) on page 771
[configure lldp ports vendor-specific avaya-extreme dot1q-framing](#) on page 772
[configure lldp ports vendor-specific avaya-extreme file-server](#) on page 774
[configure lldp ports vendor-specific avaya-extreme poe-conservation-request](#) on page 775
[configure lldp ports vendor-specific dcbx](#) on page 776
[configure lldp ports vendor-specific dot1 port-protocol-vlan-ID](#) on page 777
[configure lldp ports vendor-specific dot1 port-vlan-ID](#) on page 778
[configure lldp ports vendor-specific dot1 vlan-name](#) on page 779
[configure lldp ports vendor-specific dot3 link-aggregation](#) on page 780
[configure lldp ports vendor-specific dot3 mac-phy](#) on page 781
[configure lldp ports vendor-specific dot3 max-frame-size](#) on page 782
[configure lldp ports vendor-specific dot3 power-via-mdi](#) on page 783
[configure lldp ports vendor-specific med capabilities](#) on page 784
[configure lldp ports vendor-specific med location-identification](#) on page 786
[configure lldp ports vendor-specific med policy application](#) on page 787
[configure lldp ports vendor-specific med power-via-mdi](#) on page 789
[configure lldp reinitialize-delay](#) on page 791
[configure lldp snmp-notification-interval](#) on page 792
[configure lldp transmit-delay](#) on page 792
[configure lldp transmit-hold](#) on page 793
[configure lldp transmit-interval](#) on page 794
[configure log display](#) on page 795
[configure log filter events](#) on page 796
[configure log filter events match](#) on page 799
[configure log messages privilege](#) on page 803
[configure log target filter](#) on page 803
[configure log target format](#) on page 806
[configure log target match](#) on page 810
[configure log target memory-buffer alert percent-full](#) on page 812
[configure log target severity](#) on page 812

[configure log target syslog](#) on page 814
[configure log target upm filter](#) on page 816
[configure log target upm match](#) on page 817
[configure log target xml-notification filter](#) on page 818
[configure mac-lockdown-timeout ports aging-time](#) on page 819
[configure mac-locking ports first-arrival aging](#) on page 820
[configure mac-locking ports first-arrival limit-learning](#) on page 820
[configure mac-locking ports first-arrival link-down-action](#) on page 821
[configure mac-locking ports first-arrival move-to-static](#) on page 822
[configure mac-locking ports learn-limit-action](#) on page 823
[configure mac-locking ports log](#) on page 824
[configure mac-locking ports static delete station](#) on page 825
[configure mac-locking ports static limit-learning](#) on page 826
[configure mac-locking ports static](#) on page 827
[configure mac-locking ports trap](#) on page 828
[configure macsec cipher-suite](#) on page 829
[configure macsec connectivity-association](#) on page 830
[configure macsec include-sci](#) on page 832
[configure macsec initialize ports](#) on page 834
[configure macsec mka actor-priority](#) on page 835
[configure macsec mka life-time](#) on page 836
[configure macsec replay-protect](#) on page 838
[configure mcast ipv4 cache timeout](#) on page 839
[configure mcast ipv6 cache timeout](#) on page 840
[configure meter](#) on page 841
[configure mirror add](#) on page 844
[configure mirror add ports anomaly](#) on page 845
[configure mirror control_index](#) on page 846
[configure mirror delete](#) on page 847
[configure mirror description](#) on page 848
[configure mirror name](#) on page 849
[configure mirror to](#) on page 849
[configure mirror to remote-ip delete](#) on page 852
[configure mirror to remote-ip protocol-type](#) on page 853
[configure mlag peer alternate ipaddress](#) on page 854
[configure mlag peer authentication](#) on page 856
[configure mlag peer interval](#) on page 858
[configure mlag peer ipaddress](#) on page 858
[configure mlag peer lacp-mac](#) on page 859
[configure mlag peer name](#) on page 860
[configure mlag ports convergence-control](#) on page 861
[configure mlag ports link-up-isolation](#) on page 862

[configure mlag ports reload-delay](#) on page 863
[configure mlag ports reload-interval](#) on page 864
[configure mld](#) on page 865
[configure mld snooping fast-learning](#) on page 866
[configure mld snooping filters](#) on page 867
[configure mld snooping flood-list](#) on page 868
[configure mld snooping leave-timeout](#) on page 870
[configure mld snooping timer](#) on page 871
[configure mld snooping vlan ports add dynamic group](#) on page 872
[configure mld snooping vlan ports add static group](#) on page 873
[configure mld snooping vlan ports add static router](#) on page 874
[configure mld snooping vlan ports delete static group](#) on page 875
[configure mld snooping vlan ports delete static router](#) on page 876
[configure mld snooping vlan ports filter](#) on page 877
[configure mld snooping vlan ports join-limit](#) on page 879
[configure mld ssm-map add](#) on page 879
[configure mld ssm-map delete](#) on page 881
[configure mpls add vlan](#) on page 882
[configure mpls delete vlan](#) on page 883
[configure mpls exp examination](#) on page 883
[configure mpls exp replacement](#) on page 884
[configure mpls labels max-static](#) on page 885
[configure mpls ldp advertise](#) on page 887
[configure mpls ldp loop-detection](#) on page 888
[configure mpls ldp timers](#) on page 889
[configure mpls lsr-id](#) on page 891
[configure mpls rsvp-te bandwidth committed-rate](#) on page 892
[configure mpls rsvp-te lsp add path](#) on page 893
[configure mpls rsvp-te lsp change](#) on page 895
[configure mpls rsvp-te lsp delete path](#) on page 896
[configure mpls rsvp-te lsp fast-reroute](#) on page 896
[configure mpls rsvp-te lsp path use profile](#) on page 897
[configure mpls rsvp-te lsp transport](#) on page 898
[configure mpls rsvp-te metric](#) on page 900
[configure mpls rsvp-te path add ero](#) on page 900
[configure mpls rsvp-te path delete ero](#) on page 902
[configure mpls rsvp-te profile \(fast-reroute\)](#) on page 903
[configure mpls rsvp-te profile](#) on page 905
[configure mpls rsvp-te timers lsp rapid-retry](#) on page 908
[configure mpls rsvp-te timers lsp standard-retry](#) on page 909
[configure mpls rsvp-te timers session](#) on page 911
[configure mpls static lsp transport](#) on page 913

[configure mpls static lsp](#) on page 914
[configure mrp ports timers](#) on page 915
[configure msdp as-display-format](#) on page 917
[configure msdp max-rejected-cache](#) on page 917
[configure msdp originator-id](#) on page 919
[configure msdp peer default-peer](#) on page 920
[configure msdp peer description](#) on page 921
[configure msdp peer mesh-group](#) on page 922
[configure msdp peer no-default-peer](#) on page 923
[configure msdp peer password](#) on page 924
[configure msdp peer sa-filter](#) on page 925
[configure msdp peer sa-limit](#) on page 926
[configure msdp peer source-interface](#) on page 927
[configure msdp peer timer](#) on page 929
[configure msdp peer ttl-threshold](#) on page 930
[configure msdp sa-cache-server](#) on page 931
[configure msrp latency-max-frame-size](#) on page 932
[configure msrp ports sr-pvid](#) on page 933
[configure msrp ports traffic-class delta-bandwidth](#) on page 934
[configure msrp sharing](#) on page 935
[configure msrp timers first-value-change-recovery](#) on page 936
[configure mstp format](#) on page 937
[configure mstp region](#) on page 938
[configure mstp revision](#) on page 939
[configure mvr add receiver](#) on page 940
[configure mvr add vlan](#) on page 941
[configure mvr delete receiver](#) on page 942
[configure mvr delete vlan](#) on page 943
[configure mvr mvr-address](#) on page 944
[configure mvr static group](#) on page 945
[configure mvrp stpd](#) on page 946
[configure mvrp tag ports registration](#) on page 947
[configure mvrp tag ports transmit](#) on page 948
[configure mvrp vlan auto-creation](#) on page 949
[configure mvrp vlan registration](#) on page 950
[configure neighbor-discovery cache add](#) on page 951
[configure neighbor-discovery cache delete](#) on page 951
[configure neighbor-discovery cache locktime](#) on page 952
[configure neighbor-discovery cache max_entries](#) on page 953
[configure neighbor-discovery cache max_pending_entries](#) on page 954
[configure neighbor-discovery cache reachable-time](#) on page 955
[configure neighbor-discovery cache retransmit-time](#) on page 956

[configure neighbor-discovery cache timeout](#) on page 956

[configure netlogin add mac-list](#) on page 957

[configure netlogin add proxy-port](#) on page 959

[configure netlogin agingtime](#) on page 959

[configure netlogin allowed-refresh-failures](#) on page 960

[configure netlogin authentication database-order](#) on page 961

[configure netlogin authentication failure vlan](#) on page 962

[configure netlogin authentication protocol-order](#) on page 963

[configure netlogin authentication service-unavailable vlan](#) on page 964

[configure netlogin banner](#) on page 966

[configure netlogin base-url](#) on page 967

[configure netlogin delete mac-list](#) on page 967

[configure netlogin delete proxy-port](#) on page 968

[configure netlogin dot1x eapol-transmit-version](#) on page 969

[configure netlogin dot1x guest-vlan](#) on page 970

[configure netlogin dot1x tag-eapol](#) on page 972

[configure netlogin dot1x timers](#) on page 972

[configure netlogin dynamic-vlan](#) on page 974

[configure netlogin dynamic-vlan uplink-ports](#) on page 976

[configure netlogin idle-timeout](#) on page 978

[configure netlogin local-user security-profile](#) on page 978

[configure netlogin local-user](#) on page 979

[configure netlogin mac timers reauth-period](#) on page 981

[configure netlogin mac username case](#) on page 982

[configure netlogin mac username format](#) on page 983

[configure netlogin move-fail-action](#) on page 984

[configure netlogin port allow egress-traffic](#) on page 985

[configure netlogin ports](#) on page 986

[configure netlogin ports mode](#) on page 988

[configure netlogin ports no-restart](#) on page 991

[configure netlogin ports restart](#) on page 992

[configure netlogin redirect-page](#) on page 993

[configure netlogin session-refresh](#) on page 994

[configure netlogin session-timeout](#) on page 995

[configure netlogin trap](#) on page 996

[configure netlogin vlan](#) on page 996

[configure network-clock gtp bmca](#) on page 997

[configure network-clock gtp default-set](#) on page 998

[configure network-clock gtp ports announce](#) on page 999

[configure network-clock gtp ports peer-delay](#) on page 1000

[configure network-clock gtp ports sync](#) on page 1003

[configure network-clock gtp slave-port](#) on page 1004

[configure nodealias ports](#) on page 1005
[configure ntp key trusted/not-trusted](#) on page 1006
[configure ntp local-clock none](#) on page 1007
[configure ntp local-clock stratum](#) on page 1008
[configure ntp restrict-list](#) on page 1009
[configure ntp server/peer add](#) on page 1010
[configure ntp server/peer delete](#) on page 1011
[configure ospf bfd](#) on page 1011
[configure ospf add virtual-link](#) on page 1012
[configure ospf add vlan area](#) on page 1013
[configure ospf add vlan area link-type](#) on page 1014
[configure ospf area add range](#) on page 1015
[configure ospf area delete range](#) on page 1016
[configure ospf area external-filter](#) on page 1017
[configure ospf area interarea-filter](#) on page 1018
[configure ospf area normal](#) on page 1018
[configure ospf area nssa stub-default-cost](#) on page 1019
[configure ospf area stub stub-default-cost](#) on page 1020
[configure ospf area timer](#) on page 1021
[configure ospf ase-limit](#) on page 1023
[configure ospf ase-summary add](#) on page 1024
[configure ospf ase-summary delete](#) on page 1025
[configure ospf authentication](#) on page 1025
[configure ospf cost](#) on page 1027
[configure ospf delete virtual-link](#) on page 1028
[configure ospf delete vlan](#) on page 1028
[configure ospf import-policy](#) on page 1029
[configure ospf instanceid](#) on page 1030
[configure ospf lsa-batch-interval](#) on page 1031
[configure ospf metric-table](#) on page 1032
[configure ospf priority](#) on page 1033
[configure ospf restart grace-period](#) on page 1034
[configure ospf restart](#) on page 1035
[configure ospf restart-helper](#) on page 1036
[configure ospf routerid](#) on page 1037
[configure ospf spf-hold-time](#) on page 1038
[configure ospf virtual-link timer](#) on page 1038
[configure ospf vlan area](#) on page 1040
[configure ospf vlan neighbor add](#) on page 1041
[configure ospf vlan neighbor delete](#) on page 1041
[configure ospf vlan timer](#) on page 1042
[configure ospfv3 add interface all](#) on page 1044

[configure ospfv3 add interface](#) on page 1045
[configure ospfv3 add virtual-link](#) on page 1046
[configure ospfv3 area add range](#) on page 1047
[configure ospfv3 area cost](#) on page 1048
[configure ospfv3 area delete range](#) on page 1049
[configure ospfv3 area external-filter](#) on page 1050
[configure ospfv3 area interarea-filter](#) on page 1051
[configure ospfv3 area normal](#) on page 1052
[configure ospfv3 area nssa](#) on page 1053
[configure ospfv3 area priority](#) on page 1054
[configure ospfv3 area stub](#) on page 1055
[configure ospfv3 area timer](#) on page 1056
[configure ospfv3 authentication \(Authentication Trailer\)](#) on page 1057
[configure ospfv3 authentication \(IPsec\)](#) on page 1058
[configure ospfv3 bfd](#) on page 1060
[configure ospfv3 delete interface](#) on page 1061
[configure ospfv3 delete virtual-link](#) on page 1061
[configure ospfv3 import-policy](#) on page 1062
[configure ospfv3 interface area](#) on page 1064
[configure ospfv3 interface cost](#) on page 1064
[configure ospfv3 interface priority](#) on page 1065
[configure ospfv3 interface timer](#) on page 1066
[configure ospfv3 lsa-batch-interval](#) on page 1068
[configure ospfv3 metric-table](#) on page 1069
[configure ospfv3 restart](#) on page 1071
[configure ospfv3 restart grace-period](#) on page 1072
[configure ospfv3 restart-helper](#) on page 1072
[configure ospfv3 routerid](#) on page 1074
[configure ospfv3 spf-hold-time](#) on page 1075
[configure ospfv3 virtual-link authentication \(Authentication Trailer\)](#) on page 1076
[configure ospfv3 virtual-link authentication](#) on page 1077
[configure ospfv3 virtual-link restart-helper](#) on page 1078
[configure ospfv3 virtual-link timer](#) on page 1080
[configure pim add tunnel](#) on page 1081
[configure pim add vlan](#) on page 1082
[configure pim anycast-rp](#) on page 1084
[configure pim border](#) on page 1085
[configure pim cbsr](#) on page 1086
[configure pim crp static](#) on page 1087
[configure pim crp timer](#) on page 1088
[configure pim crp vlan](#) on page 1089
[configure pim delete tunnel](#) on page 1090

[configure pim delete vlan](#) on page 1091

[configure pim dense-neighbor-check](#) on page 1092

[configure pim dr-priority](#) on page 1093

[configure pim iproute sharing hash](#) on page 1094

[configure pim register-policy](#) on page 1095

[configure pim register-policy rp](#) on page 1096

[configure pim register-rate-limit-interval](#) on page 1097

[configure pim register-suppress-interval register-probe-interval](#) on page 1098

[configure pim snooping sgrpt-prune](#) on page 1099

[configure pim shutdown-priority](#) on page 1099

[configure pim spt-threshold](#) on page 1100

[configure pim ssm range](#) on page 1101

[configure pim state-refresh timer origination-interval](#) on page 1103

[configure pim state-refresh timer source-active-timer](#) on page 1104

[configure pim state-refresh ttl](#) on page 1104

[configure pim state-refresh](#) on page 1105

[configure pim timer tunnel](#) on page 1106

[configure pim timer vlan](#) on page 1107

[configure pim tunnel dr-priority](#) on page 1108

[configure pim vlan trusted-gateway](#) on page 1109

[configure policy access-list](#) on page 1110

[configure policy autoclear](#) on page 1111

[configure policy app-signature group name pattern](#) on page 1113

[configure policy app-signature minimum-ttl](#) on page 1114

[configure policy captive-portal](#) on page 1115

[configure policy captive-portal listening](#) on page 1116

[configure policy captive-portal rule-use](#) on page 1117

[configure policy convergence-endpoint](#) on page 1118

[configure policy convergence-endpoint clear](#) on page 1118

[configure policy convergence-endpoint index](#) on page 1119

[configure policy convergence-endpoint ports](#) on page 1120

[configure policy invalid action](#) on page 1121

[configure policy mactable](#) on page 1122

[configure policy port](#) on page 1123

[configure policy profile](#) on page 1123

[configure policy resource-profile](#) on page 1126

[configure policy rule](#) on page 1127

[configure policy rule admin-profile](#) on page 1131

[configure policy rule-model](#) on page 1132

[configure policy slices shared](#) on page 1133

[configure policy slices tci-overwrite](#) on page 1134

[configure policy syslog](#) on page 1135

[configure policy vlanauthorization](#) on page 1136
[configure policy vlanauthorization port](#) on page 1137
[configure port description-string](#) on page 1138
[configure port ethertype](#) on page 1139
[configure port reflective-relay](#) on page 1140
[configure port shared-packet-buffer](#) on page 1140
[configure ports](#) on page 1141
[configure ports auto off](#) on page 1142
[configure ports auto on](#) on page 1144
[configure ports auto-polarity](#) on page 1146
[configure ports ddmi](#) on page 1147
[configure ports display-string](#) on page 1148
[configure ports dot1p](#) on page 1149
[configure ports dwdm channel none](#) on page 1149
[configure ports dwdm channel](#) on page 1150
[configure ports eee](#) on page 1152
[configure ports forward-error-correction](#) on page 1153
[configure ports ingress-filtering](#) on page 1154
[configure ports isolation](#) on page 1155
[configure ports l2pt profile](#) on page 1156
[configure ports link-flap-detection action](#) on page 1157
[configure ports link-flap-detection interval threshold disable-time](#) on page 1158
[configure ports link-flap-detection](#) on page 1160
[configure ports link-scan interval](#) on page 1161
[configure ports monitor vlan](#) on page 1162
[configure ports partition](#) on page 1163
[configure ports protocol filter](#) on page 1166
[configure ports qosprofile](#) on page 1167
[configure ports rate-limit egress](#) on page 1167
[configure ports rate-limit flood](#) on page 1169
[configure ports redundant](#) on page 1170
[configure ports vlan](#) on page 1171
[configure power monitor](#) on page 1174
[configure private-vlan add network](#) on page 1175
[configure private-vlan add subscriber](#) on page 1176
[configure private-vlan delete](#) on page 1177
[configure protocol add](#) on page 1178
[configure process group other cpu-limit](#) on page 1179
[configure process group other memory-limit](#) on page 1180
[configure protocol delete](#) on page 1182
[configure protocol filter](#) on page 1183
[configure qosprofile](#) on page 1185

[configure qosprofile weight](#) on page 1188
[configure qosprofile wred](#) on page 1189
[configure qosprofile egress wred ecn](#) on page 1191
[configure qosscheduler weighted-deficit-round-robin](#) on page 1192
[configure radius algorithm](#) on page 1193
[configure radius retries](#) on page 1194
[configure radius server client-ip](#) on page 1195
[configure radius shared-secret](#) on page 1197
[configure radius timeout](#) on page 1198
[configure radius-accounting retries](#) on page 1200
[configure radius-accounting server client-ip](#) on page 1200
[configure radius-accounting shared-secret](#) on page 1202
[configure radius-accounting timeout](#) on page 1204
[configure radius dynamic-authorization server client-ip](#) on page 1205
[configure radius tls ocsf](#) on page 1207
[configure radius tls ocsf nonce](#) on page 1207
[configure radius tls ocsf override](#) on page 1208
[configure radius tls ocsf signer](#) on page 1209
[configure radius tls tcp-user-timeout](#) on page 1210
[configure rip add vlan](#) on page 1211
[configure rip delete vlan](#) on page 1212
[configure rip garbage-time](#) on page 1213
[configure rip import-policy](#) on page 1213
[configure rip route-timeout](#) on page 1214
[configure rip update-time](#) on page 1215
[configure rip vlan cost](#) on page 1216
[configure rip vlan route-policy](#) on page 1216
[configure rip vlan rx-mode](#) on page 1218
[configure rip vlan trusted-gateway](#) on page 1218
[configure rip vlan tx-mode](#) on page 1219
[configure ripng add](#) on page 1220
[configure ripng cost](#) on page 1221
[configure ripng delete](#) on page 1222
[configure ripng garbage-time](#) on page 1223
[configure ripng import-policy](#) on page 1224
[configure ripng route-policy](#) on page 1225
[configure ripng route-timeout](#) on page 1226
[configure ripng trusted-gateway](#) on page 1227
[configure ripng update-time](#) on page 1228
[configure switch safe-default-script](#) on page 1229
[configure security fips-mode](#) on page 1230
[configure security python](#) on page 1231

[configure sflow agent ipaddress](#) on page 1232
[configure sflow collector ipaddress](#) on page 1233
[configure sflow max-cpu-sample-limit](#) on page 1234
[configure sflow poll-interval](#) on page 1235
[configure sflow ports sample-rate](#) on page 1236
[configure sflow sample-rate](#) on page 1236
[configure sharing add ports](#) on page 1237
[configure sharing address-based custom](#) on page 1239
[configure sharing address-based custom hash-seed](#) on page 1240
[configure sharing algorithm](#) on page 1241
[configure sharing delete ports](#) on page 1243
[configure sharing distribution-mode](#) on page 1244
[configure sharing health-check member-port add tcp-tracking](#) on page 1245
[configure sharing health-check member-port delete tcp-tracking](#) on page 1246
[configure sharing health-check member-port tcp-tracking](#) on page 1247
[configure sharing lacp activity-mode](#) on page 1248
[configure sharing lacp defaulted-state-action](#) on page 1249
[configure sharing lacp fallback](#) on page 1250
[configure sharing lacp fallback timeout](#) on page 1252
[configure sharing lacp system-priority](#) on page 1253
[configure sharing lacp timeout](#) on page 1254
[configure sharing minimum-active](#) on page 1255
[configure sharing port-based key](#) on page 1256
[configure slot description](#) on page 1257
[configure slot module](#) on page 1258
[configure slot restart-limit](#) on page 1259
[configure slpp guard ethertype](#) on page 1260
[configure slpp guard recovery-timeout](#) on page 1261
[configure snmp access-profile](#) on page 1262
[configure snmp add community](#) on page 1264
[configure snmp add notification-log](#) on page 1265
[configure snmp add trapreceiver](#) on page 1266
[configure snmp delete community](#) on page 1268
[configure snmp delete notification-log](#) on page 1269
[configure snmp delete trapreceiver](#) on page 1270
[configure snmp ifmibifalias size](#) on page 1271
[configure snmp notification-log filter-profile-name](#) on page 1272
[configure snmp notification-log](#) on page 1273
[configure snmp sysContact](#) on page 1274
[configure snmp sysLocation](#) on page 1275
[configure snmp sysName](#) on page 1276
[configure snmp traps batch-delay bfd](#) on page 1277

[configure snmpv3 add access](#) on page 1278
[configure snmpv3 add community](#) on page 1280
[configure snmpv3 add filter](#) on page 1282
[configure snmpv3 add filter-profile](#) on page 1283
[configure snmpv3 add group user](#) on page 1284
[configure snmpv3 add mib-view](#) on page 1286
[configure snmpv3 add notify](#) on page 1287
[configure snmpv3 add target-addr](#) on page 1288
[configure snmpv3 add target-params](#) on page 1290
[configure snmpv3 add user](#) on page 1292
[configure snmpv3 add user clone-from](#) on page 1294
[configure snmpv3 delete access](#) on page 1295
[configure snmpv3 delete community](#) on page 1296
[configure snmpv3 delete filter](#) on page 1298
[configure snmpv3 delete filter-profile](#) on page 1299
[configure snmpv3 delete group user](#) on page 1300
[configure snmpv3 delete mib-view](#) on page 1301
[configure snmpv3 delete notify](#) on page 1302
[configure snmpv3 delete target-addr](#) on page 1303
[configure snmpv3 delete target-params](#) on page 1304
[configure snmpv3 delete user](#) on page 1305
[configure snmpv3 engine-boots](#) on page 1306
[configure snmpv3 engine-id](#) on page 1307
[configure snmpv3 target-addr retry](#) on page 1308
[configure snmpv3 target-addr timeout](#) on page 1308
[configure snmp-client](#) on page 1309
[configure snmp-client update-interval](#) on page 1310
[configure ssh2 access-profile](#) on page 1311
[configure ssh2 dh-group](#) on page 1313
[configure ssh2 disable cipher mac](#) on page 1314
[configure ssh2 disable pk-alg](#) on page 1315
[configure ssh2 enable cipher mac](#) on page 1316
[configure ssh2 enable pk-alg](#) on page 1317
[configure ssh2 idletimeout](#) on page 1317
[configure ssh2 key](#) on page 1319
[configure ssh2 login-grace-timeout](#) on page 1320
[configure ssh2 rekey](#) on page 1321
[configure ssh2 secure-mode](#) on page 1322
[configure ssh2 x509v3 ocsp](#) on page 1324
[configure ssh2 x509v3 ocsp nonce](#) on page 1325
[configure ssh2 x509v3 ocsp override](#) on page 1325
[configure ssh2 x509v3 ocsp signer](#) on page 1326

[configure ssh2 x509v3 radius-password-auth](#) on page 1327
[configure ssh2 x509v3 username overwrite](#) on page 1328
[configure ssh2 x509v3 username strip-domain](#) on page 1329
[configure ssh2 x509v3 username use-domain](#) on page 1330
[configure sshd2 user-key add user](#) on page 1331
[configure sshd2 user-key delete user](#) on page 1332
[configure ssl certificate hash-algorithm](#) on page 1333
[configure ssl certificate pregenerated](#) on page 1334
[configure ssl certificate privkeylen](#) on page 1335
[configure ssl csr](#) on page 1336
[configure ssl privkey pregenerated](#) on page 1338
[configure stack-ports debounce time](#) on page 1339
[configure stacking alternate-ip-address](#) on page 1340
[configure stacking easy-setup](#) on page 1342
[configure stacking license-level](#) on page 1344
[configure stacking mac-address](#) on page 1346
[configure stacking master-capability](#) on page 1348
[configure stacking node-address](#) on page 1349
[configure stacking priority](#) on page 1350
[configure stacking redundancy](#) on page 1351
[configure stacking slot-number automatic](#) on page 1353
[configure stacking-support auto-discovery](#) on page 1354
[configure stacking-support stack-ports](#) on page 1355
[configure stpd add vlan](#) on page 1358
[configure stpd backup-root](#) on page 1361
[configure stpd bpdu-forwarding](#) on page 1362
[configure stpd default-encapsulation](#) on page 1363
[configure stpd delete vlan](#) on page 1365
[configure stpd description](#) on page 1366
[configure stpd filter-method](#) on page 1367
[configure stpd flush-method](#) on page 1368
[configure stpd forwarddelay](#) on page 1369
[configure stpd hellotime](#) on page 1370
[configure stpd loop-protect event-threshold](#) on page 1371
[configure stpd loop-protect event-window](#) on page 1371
[configure stpd maxage](#) on page 1372
[configure stpd max-hop-count](#) on page 1373
[configure stpd mode](#) on page 1374
[configure stpd multicast send-query](#) on page 1376
[configure stpd ports active-role disable](#) on page 1377
[configure stpd ports active-role enable](#) on page 1378
[configure stpd ports auto-edge](#) on page 1379

[configure stpd ports bpdu-restrict](#) on page 1380
[configure stpd ports cost](#) on page 1380
[configure stpd ports edge-safeguard disable](#) on page 1382
[configure stpd ports edge-safeguard enable](#) on page 1384
[configure stpd ports link-type](#) on page 1386
[configure stpd ports loop-protect](#) on page 1388
[configure stpd ports loop-protect partner](#) on page 1389
[configure stpd ports mode](#) on page 1390
[configure stpd ports port-priority](#) on page 1391
[configure stpd ports priority](#) on page 1393
[configure stpd ports reflection-bpdu](#) on page 1394
[configure stpd ports restricted-role disable](#) on page 1395
[configure stpd ports restricted-role enable](#) on page 1396
[configure stpd ports restricted-tcn](#) on page 1397
[configure stpd priority](#) on page 1398
[configure stpd priority-mode](#) on page 1399
[configure stpd tag](#) on page 1400
[configure stpd trap new-root](#) on page 1402
[configure stpd trap topology-change](#) on page 1402
[configure stpd tx-hold-count](#) on page 1403
[configure switch integrity-check image](#) on page 1404
[configure sys-health-check all level](#) on page 1405
[configure syslog add](#) on page 1407
[configure syslog tls cipher](#) on page 1408
[configure syslog tls ocsp](#) on page 1410
[configure syslog tls ocsp nonce](#) on page 1410
[configure syslog tls ocsp override](#) on page 1411
[configure syslog tls ocsp signer](#) on page 1412
[configure syslog tls tcp-user-timeout](#) on page 1413
[configure syslog delete](#) on page 1414
[configure syslog reference-identifier](#) on page 1416
[configure system ports notation](#) on page 1417
[configure sys-recovery-level switch](#) on page 1418
[configure sys-recovery-level](#) on page 1420
[configure tacacs priv-lvl](#) on page 1421
[configure tacacs server client-ip](#) on page 1422
[configure tacacs shared-secret](#) on page 1423
[configure tacacs timeout](#) on page 1424
[configure tacacs-accounting server](#) on page 1425
[configure tacacs-accounting shared-secret](#) on page 1426
[configure tacacs-accounting timeout](#) on page 1427
[configure tech-support add collector](#) on page 1428

[configure tech-support collector](#) on page 1430
[configure tech-support collector data-set](#) on page 1431
[configure tech-support collector frequency error-detected](#) on page 1432
[configure tech-support collector report](#) on page 1433
[configure tech-support delete collector](#) on page 1434
[configure telnet access-profile](#) on page 1435
[configure telnet port](#) on page 1438
[configure telnet vr](#) on page 1439
[configure time](#) on page 1440
[configure timezone](#) on page 1441
[configure trusted-ports trust-for dhcp-server](#) on page 1444
[configure trusted-servers add server](#) on page 1446
[configure trusted-servers delete server](#) on page 1447
[configure tunnel ipaddress](#) on page 1448
[configure tunnel ip tcp adjust-mss](#) on page 1449
[configure tunnel ipv6 tcp adjust-mss](#) on page 1450
[configure twamp endpoint](#) on page 1451
[configure twamp key-id](#) on page 1452
[configure twamp reflector](#) on page 1452
[configure twamp server](#) on page 1453
[configure upm event](#) on page 1454
[configure upm profile maximum execution-time](#) on page 1455
[configure upm timer after](#) on page 1456
[configure upm timer at](#) on page 1457
[configure upm timer profile](#) on page 1458
[configure virtual-network](#) on page 1458
[configure virtual-network add network ports](#) on page 1460
[configure virtual-network delete network ports](#) on page 1460
[configure virtual-network dynamic](#) on page 1461
[configure virtual-network flooding](#) on page 1462
[configure virtual-network local endpoint](#) on page 1463
[configure virtual-network monitor](#) on page 1465
[configure virtual-network multicast group](#) on page 1466
[configure virtual-network name](#) on page 1466
[configure virtual-network remote-endpoint vxlan ipaddress](#) on page 1467
[configure virtual-network remote-endpoint vxlan ipaddress monitor](#) on page 1468
[configure virtual-network replication-role](#) on page 1469
[configure virtual-network selected-replicator](#) on page 1471
[configure virtual-network vxlan vni](#) on page 1472
[configure vlan add nsi | isid](#) on page 1472
[configure vlan add ports](#) on page 1473
[configure vlan add ports private-vlan translated](#) on page 1475

[configure vlan add ports stpd](#) on page 1476
[configure vlan add secondary-ipaddress](#) on page 1479
[configure vlan delete nsi | isid](#) on page 1480
[configure vlan delete ports](#) on page 1481
[configure vlan delete secondary-ipaddress](#) on page 1482
[configure vlan description](#) on page 1483
[configure vlan dhcp-address-range](#) on page 1484
[configure vlan dhcp-lease-timer](#) on page 1485
[configure vlan dhcp-options](#) on page 1486
[configure vlan dynamic-vlan uplink-ports](#) on page 1487
[configure vlan ipaddress](#) on page 1488
[configure vlan l2pt profile](#) on page 1490
[configure vlan name](#) on page 1491
[configure vlan netlogin-lease-timer](#) on page 1492
[configure vlan qosprofile](#) on page 1493
[configure vlan protocol](#) on page 1494
[configure vlan router-discovery add prefix](#) on page 1495
[configure vlan router-discovery default-lifetime](#) on page 1496
[configure vlan router-discovery delete prefix](#) on page 1497
[configure vlan router-discovery link-mtu](#) on page 1498
[configure vlan router-discovery managed-config-flag](#) on page 1499
[configure vlan router-discovery max-interval](#) on page 1500
[configure vlan router-discovery min-interval](#) on page 1500
[configure vlan router-discovery other-config-flag](#) on page 1501
[configure vlan router-discovery reachable-time](#) on page 1502
[configure vlan router-discovery retransmit-time](#) on page 1503
[configure vlan router-discovery set prefix](#) on page 1504
[configure router-discovery vrrp-lla-only](#) on page 1505
[configure vlan subvlan](#) on page 1506
[configure vlan subvlan-address-range](#) on page 1507
[configure vlan suppress](#) on page 1508
[configure vlan tag](#) on page 1509
[configure vlan udp-profile](#) on page 1510
[configure vlan untagged-ports auto-move](#) on page 1512
[configure vlan-translation add loopback-port](#) on page 1514
[configure vlan-translation add member-vlan](#) on page 1514
[configure vlan-translation delete loopback-port](#) on page 1515
[configure vlan-translation delete member-vlan](#) on page 1516
[configure vm add | delete ports](#) on page 1517
[configure vm add virtual-interface](#) on page 1518
[configure vm delete virtual-interface](#) on page 1519
[configure vm cpus](#) on page 1520

[configure vm disk bus-type](#) on page 1521
[configure vm memory](#) on page 1522
[configure vm vnc](#) on page 1523
[configure vman add ports](#) on page 1524
[configure vman add ports cep](#) on page 1526
[configure vman delete ports](#) on page 1528
[configure vman ethertype](#) on page 1529
[configure vman ports add cvid](#) on page 1530
[configure vman ports delete cvid](#) on page 1531
[configure vman protocol](#) on page 1532
[configure vman tag](#) on page 1534
[configure vm-tracking authentication database-order](#) on page 1534
[configure vm-tracking blackhole](#) on page 1535
[configure vm-tracking local-vm](#) on page 1536
[configure vm-tracking nms timeout](#) on page 1537
[configure vm-tracking nms](#) on page 1538
[configure vm-tracking repository](#) on page 1539
[configure vm-tracking timers](#) on page 1540
[configure vm-tracking vpp add](#) on page 1541
[configure vm-tracking vpp counters](#) on page 1542
[configure vm-tracking vpp delete](#) on page 1543
[configure vm-tracking vpp vlan-tag](#) on page 1544
[configure vpex auto-configuration mlag-id](#) on page 1545
[configure vpex mlag-id peer](#) on page 1546
[configure vpex ports](#) on page 1547
[configure vpex ring rebalancing](#) on page 1548
[configure vpls](#) on page 1550
[configure vpls add peer](#) on page 1552
[configure vpls delete peer](#) on page 1553
[configure vpls delete service](#) on page 1554
[configure vpls health-check vccv](#) on page 1555
[configure vpls peer l2pt profile](#) on page 1556
[configure vpls peer mpls lsp](#) on page 1558
[configure vpls peer](#) on page 1559
[configure vpls snmp-vpn-identifier](#) on page 1560
[configure vr add ports](#) on page 1561
[configure vr add protocol](#) on page 1562
[configure vr delete ports](#) on page 1564
[configure vr description](#) on page 1564
[configure vrrp group](#) on page 1565
[configure vrrp fabric-routing](#) on page 1567
[configure vrrp vlan vrid accept-mode](#) on page 1568

[configure vrrp vlan vrid add ipaddress](#) on page 1569
[configure vrrp vlan vrid add track-iproute](#) on page 1571
[configure vrrp vlan vrid add track-ping](#) on page 1572
[configure vrrp vlan vrid add track-vlan](#) on page 1573
[configure vrrp vlan vrid add virtual-link-local](#) on page 1574
[configure vrrp vlan vrid advertisement-interval](#) on page 1575
[configure vrrp vlan vrid delete track-iproute](#) on page 1576
[configure vrrp vlan vrid delete track-ping](#) on page 1577
[configure vrrp vlan vrid delete track-vlan](#) on page 1578
[configure vrrp vlan vrid delete ipaddress](#) on page 1579
[configure vrrp vlan vrid dont-preempt](#) on page 1580
[configure vrrp vlan vrid host-mobility](#) on page 1581
[configure vrrp vlan vrid ipv4 checksum](#) on page 1582
[configure vrrp vlan vrid preempt](#) on page 1583
[configure vrrp vlan vrid priority](#) on page 1584
[configure vrrp vlan vrid track-mode](#) on page 1585
[configure vrrp vlan vrid version](#) on page 1586
[configure web http access-profile](#) on page 1587
[configure xml-notification target add/delete](#) on page 1589
[configure xml-notification target](#) on page 1590
[configure l2pt encapsulation dest-mac](#) on page 1591
[cp](#) on page 1591
[create access-list](#) on page 1594
[create access-list network-zone](#) on page 1596
[create access-list zone](#) on page 1597
[create account](#) on page 1597
[create auto-peering bgp](#) on page 1600
[create auto-peering ospf](#) on page 1601
[create bgp evpn instance](#) on page 1602
[create bgp neighbor peer-group](#) on page 1602
[create bgp neighbor remote-AS-number](#) on page 1604
[create bgp peer-group](#) on page 1606
[create cfm domain dns md-level](#) on page 1607
[create cfm domain mac md-level](#) on page 1608
[create cfm domain string md-level](#) on page 1609
[create cfm segment destination](#) on page 1610
[create database](#) on page 1612
[create eaps shared-port](#) on page 1612
[create eaps](#) on page 1613
[create erps ring](#) on page 1614
[create esrp](#) on page 1615
[create fdb mac-tracking entry](#) on page 1616

[create fdb vlan ports](#) on page 1617
[create flow-redirect](#) on page 1619
[create flowmon collector](#) on page 1620
[create flowmon group](#) on page 1621
[create flowmon key](#) on page 1622
[create identity-management role](#) on page 1623
[create ip nat rule](#) on page 1626
[create isis area](#) on page 1627
[create keychain](#) on page 1628
[create l2pt profile](#) on page 1629
[create l2vpn fec-id-type pseudo-wire](#) on page 1630
[create ldap domain](#) on page 1631
[create log filter](#) on page 1632
[create log message](#) on page 1633
[create log target upm](#) on page 1634
[create log target xml-notification](#) on page 1635
[create macsec connectivity-association](#) on page 1636
[create meter](#) on page 1637
[create mirror control_index](#) on page 1638
[create mirror](#) on page 1639
[create mlag peer](#) on page 1641
[create mpls rsvp-te path](#) on page 1642
[create mpls rsvp-te profile fast-reroute](#) on page 1644
[create mpls rsvp-te profile](#) on page 1645
[create mpls static lsp](#) on page 1646
[create msdp mesh-group](#) on page 1647
[create msdp peer](#) on page 1648
[create netlogin local-user](#) on page 1649
[create ntp key](#) on page 1651
[create ospf area](#) on page 1652
[create ospfv3 area](#) on page 1653
[create policy access-list](#) on page 1654
[create policy access-list action-set](#) on page 1657
[create ports group](#) on page 1658
[create private-vlan](#) on page 1659
[create process executable](#) on page 1660
[create process python-module](#) on page 1661
[create protocol](#) on page 1663
[create qosprofile](#) on page 1664
[create snmp trap](#) on page 1665
[create sshd2 key-file](#) on page 1665
[create sshd2 user-key](#) on page 1666

[create stpd](#) on page 1667
[create tunnel 6to4](#) on page 1668
[create tunnel gre destination source](#) on page 1669
[create tunnel ipv6-in-ipv4](#) on page 1670
[create upm profile](#) on page 1671
[create upm timer](#) on page 1672
[create virtual-network](#) on page 1673
[create virtual-network remote-endpoint vxlan ipaddress](#) on page 1675
[create virtual-router](#) on page 1675
[create vlan](#) on page 1677
[create vm image](#) on page 1680
[create vm ova](#) on page 1681
[create vman](#) on page 1683
[create vm-tracking local-vm](#) on page 1684
[create vm-tracking vpp](#) on page 1685
[create vpls fec-id-type pseudo-wire](#) on page 1686
[create vrrp group](#) on page 1687
[create vrrp vlan vrid](#) on page 1688
[create xml-notification target url](#) on page 1689
[delete access-list](#) on page 1690
[delete access-list network-zone](#) on page 1691
[delete access-list zone](#) on page 1692
[delete account](#) on page 1693
[delete auto-peering](#) on page 1694
[delete bgp evpn instance](#) on page 1695
[delete bgp neighbor](#) on page 1695
[delete bgp peer-group](#) on page 1696
[delete cfm domain](#) on page 1697
[delete cfm segment](#) on page 1698
[delete database](#) on page 1699
[delete eaps shared-port](#) on page 1699
[delete eaps](#) on page 1700
[delete erps](#) on page 1701
[delete esrp](#) on page 1702
[delete fdb mac-tracking entry](#) on page 1703
[delete fdb](#) on page 1703
[delete flow-redirect](#) on page 1705
[delete flowmon collector](#) on page 1705
[delete flowmon group](#) on page 1706
[delete flowmon key](#) on page 1707
[delete identity-management role](#) on page 1708
[delete ip nat rule](#) on page 1709

[delete isis area](#) on page 1709
[delete keychain](#) on page 1710
[delete l2pt profile](#) on page 1711
[delete l2vpn](#) on page 1712
[delete ldap domain](#) on page 1713
[delete log filter](#) on page 1714
[delete log target upm](#) on page 1715
[delete log target xml-notification](#) on page 1715
[delete macsec connectivity-association](#) on page 1716
[delete meter](#) on page 1717
[delete mirror name](#) on page 1718
[delete mlag peer](#) on page 1719
[delete mpls rsvp-te lsp](#) on page 1720
[delete mpls rsvp-te path](#) on page 1721
[delete mpls rsvp-te profile](#) on page 1721
[delete mpls static lsp](#) on page 1722
[delete msdp mesh-group](#) on page 1723
[delete msdp peer](#) on page 1724
[delete netlogin local-user](#) on page 1725
[delete ntp key](#) on page 1726
[delete ospf area](#) on page 1727
[delete ospfv3 area](#) on page 1727
[delete policy access-list](#) on page 1728
[delete policy access-list action-set](#) on page 1729
[delete ports group](#) on page 1730
[delete private-vlan](#) on page 1731
[delete process](#) on page 1732
[delete protocol](#) on page 1732
[delete qosprofile](#) on page 1733
[delete sshd2 user-key](#) on page 1734
[delete stpd](#) on page 1735
[delete tunnel](#) on page 1736
[delete upm profile](#) on page 1736
[delete upm timer](#) on page 1737
[delete var](#) on page 1738
[delete var key](#) on page 1739
[delete virtual-network](#) on page 1739
[delete virtual-network remote-endpoint vxlan ipaddress](#) on page 1740
[delete virtual-router](#) on page 1741
[delete vlan](#) on page 1742
[delete vman](#) on page 1743
[delete vm](#) on page 1744

[delete vm-tracking local-vm](#) on page 1744
[delete vm-tracking vpp](#) on page 1745
[delete vpls](#) on page 1746
[delete vrrp group](#) on page 1747
[delete vrrp vlan vrid](#) on page 1748
[delete xml-notification target](#) on page 1748
[disable access-list permit to-cpu](#) on page 1749
[disable access-list refresh blackhole](#) on page 1750
[disable account](#) on page 1751
[disable auto-provision](#) on page 1752
[disable avb](#) on page 1753
[disable avb ports](#) on page 1754
[disable bgp](#) on page 1754
[disable bgp advertise-inactive-route](#) on page 1755
[disable bgp aggregation](#) on page 1756
[disable bgp always-compare-med](#) on page 1757
[disable bgp community format](#) on page 1758
[disable bgp export vr](#) on page 1759
[disable bgp export](#) on page 1760
[disable bgp export \[static | direct\] l2vpn-evpn](#) on page 1762
[disable bgp fast-external-fallover](#) on page 1763
[disable bgp mpls-next-hop](#) on page 1764
[disable bgp multipath-relax](#) on page 1765
[disable bgp neighbor address-family l2vpn-evpn](#) on page 1766
[disable bgp neighbor capability address-family vpnv4](#) on page 1767
[disable bgp neighbor capability](#) on page 1768
[disable bgp neighbor originate-default](#) on page 1770
[disable bgp neighbor remove-private-AS-numbers](#) on page 1771
[disable bgp neighbor soft-in-reset](#) on page 1772
[disable bgp neighbor](#) on page 1774
[disable bgp peer-group capability address-family vpnv4](#) on page 1775
[disable bgp peer-group capability](#) on page 1776
[disable bgp peer-group next-hop-unchanged](#) on page 1777
[disable bgp peer-group originate-default](#) on page 1778
[disable bgp peer-group remove-private-AS-numbers](#) on page 1779
[disable bgp peer-group soft-in-reset](#) on page 1780
[disable bgp peer-group](#) on page 1781
[disable bootp vlan](#) on page 1782
[disable bootprelay ipv6](#) on page 1783
[disable bootprelay](#) on page 1784
[disable cdp ports](#) on page 1785
[disable cfm segment frame-delay measurement](#) on page 1786

[disable cfm segment frame-loss measurement mep](#) on page 1787

[disable clear-flow](#) on page 1788

[disable cli history expansion](#) on page 1788

[disable cli prompting](#) on page 1789

[disable cli refresh](#) on page 1790

[disable cli scripting](#) on page 1791

[disable cli scripting output](#) on page 1792

[disable cli space-completion](#) on page 1793

[disable cli config-logging](#) on page 1794

[disable cli-config-logging expansion](#) on page 1795

[disable cli paging](#) on page 1796

[disable cpu-monitoring](#) on page 1797

[disable dhcp ports vlan](#) on page 1797

[disable dhcp vlan](#) on page 1798

[disable diffserv examination ports](#) on page 1799

[disable diffserv replacement ports](#) on page 1800

[disable dns cache](#) on page 1801

[disable dns cache analytics](#) on page 1802

[disable dns cache dnssec](#) on page 1803

[disable dos-protect](#) on page 1803

[disable dot1p examination inner-tag ports](#) on page 1804

[disable dot1p examination ports](#) on page 1805

[disable dot1p replacement ports](#) on page 1806

[disable eaps](#) on page 1807

[disable edp ports](#) on page 1809

[disable elrp-client](#) on page 1809

[disable elsm ports](#) on page 1810

[disable elsm ports auto-restart](#) on page 1811

[disable erps](#) on page 1812

[disable erps block-vc-recovery](#) on page 1813

[disable erps ring-name](#) on page 1814

[disable erps topology-change](#) on page 1815

[disable esrp](#) on page 1815

[disable ethernet oam ports link-fault-management](#) on page 1816

[disable fdb static-mac-move](#) on page 1817

[disable flooding ports](#) on page 1818

[disable flow-control ports](#) on page 1819

[disable flowmon](#) on page 1821

[disable flowmon group](#) on page 1822

[disable icmp ipv6](#) on page 1823

[disable icmp redirects ipv6 fast-path](#) on page 1823

[disable icmp redirects](#) on page 1824

[disable icmp userredirects](#) on page 1825
[disable identity-management](#) on page 1826
[disable cli idletimeout](#) on page 1827
[disable igmp](#) on page 1828
[disable igmp snooping vlan fast-leave](#) on page 1829
[disable igmp snooping](#) on page 1830
[disable igmp ssm-map](#) on page 1831
[disable inline-power](#) on page 1832
[disable inline-power ports](#) on page 1833
[disable inline-power slot](#) on page 1834
[disable ip anycast](#) on page 1835
[disable ip nat](#) on page 1836
[disable ip nat rule](#) on page 1837
[disable iparp checking](#) on page 1838
[disable iparp gratuitous protect vlan](#) on page 1839
[disable iparp refresh](#) on page 1840
[disable ipforwarding broadcast](#) on page 1840
[disable ipforwarding broadcast](#) on page 1841
[disable ipforwarding ipv6](#) on page 1842
[disable ipmcforwarding ipv6](#) on page 1843
[disable ipmcforwarding](#) on page 1844
[disable ipmcforwarding tunnel](#) on page 1845
[disable ip option loose-source-route](#) on page 1846
[disable ip option strict-source-route](#) on page 1847
[disable iproute bfd](#) on page 1847
[disable iproute bfd strict](#) on page 1848
[disable iproute compression](#) on page 1849
[disable iproute ipv6 compression](#) on page 1850
[disable iproute ipv6 sharing](#) on page 1851
[disable iproute mpls-next-hop](#) on page 1851
[disable iproute protection ping](#) on page 1852
[disable iproute sharing](#) on page 1853
[disable ip-security anomaly-protection icmp](#) on page 1854
[disable ip-security anomaly-protection ip](#) on page 1855
[disable ip-security anomaly-protection l4port](#) on page 1855
[disable ip-security anomaly-protection notify](#) on page 1856
[disable ip-security anomaly-protection tcp flags](#) on page 1857
[disable ip-security anomaly-protection tcp fragment](#) on page 1858
[disable ip-security anomaly-protection](#) on page 1858
[disable ip-security arp gratuitous-protection](#) on page 1859
[disable ip-security arp learning learn-from-arp](#) on page 1860
[disable ip-security arp learning learn-from-dhcp](#) on page 1862

[disable ip-security arp validation](#) on page 1863
[disable ip-security dhcp-bindings restoration](#) on page 1864
[disable ip-security dhcp-snooping](#) on page 1864
[disable ip-security source-ip-lockdown ports](#) on page 1865
[disable iqagent](#) on page 1866
[disable irdp](#) on page 1867
[disable isis](#) on page 1868
[disable isis area adjacency-check](#) on page 1868
[disable isis area dynamic-hostname](#) on page 1869
[disable isis area export ipv6](#) on page 1870
[disable isis area export](#) on page 1871
[disable isis area originate-default](#) on page 1872
[disable isis area overload-bit](#) on page 1873
[disable isis hello-padding](#) on page 1874
[disable isis restart-helper](#) on page 1875
[disable jumbo-frame ports](#) on page 1875
[disable l2vpn](#) on page 1876
[disable l2vpn health-check vccv](#) on page 1877
[disable l2vpn service](#) on page 1878
[disable l2vpn sharing](#) on page 1879
[disable l2vpn vpls peer fdb send-mac-withdrawal](#) on page 1880
[disable learning iparp sender-mac](#) on page 1881
[disable learning port](#) on page 1882
[disable learning vxlan ipaddress](#) on page 1883
[disable led locator](#) on page 1883
[disable lldp ports](#) on page 1884
[disable log debug-mode](#) on page 1885
[disable log display](#) on page 1886
[disable log target](#) on page 1887
[disable log target upm](#) on page 1888
[disable log target xml-notification](#) on page 1889
[disable loopback-mode vlan](#) on page 1890
[disable mac-lockdown-timeout ports](#) on page 1891
[disable mac-locking ports](#) on page 1892
[disable mac-locking](#) on page 1892
[disable mirror](#) on page 1893
[disable mirror control_index](#) on page 1894
[disable mlag port](#) on page 1895
[disable mlag port reload-delay](#) on page 1896
[disable mld](#) on page 1896
[disable mld snooping](#) on page 1897
[disable mld-ssm map](#) on page 1898

[disable mpls](#) on page 1899

[disable mpls bfd](#) on page 1900

[disable mpls exp examination](#) on page 1901

[disable mpls exp replacement](#) on page 1902

[disable mpls ldp bgp-routes](#) on page 1902

[disable mpls ldp loop-detection](#) on page 1903

[disable mpls ldp](#) on page 1904

[disable mpls php](#) on page 1905

[disable mpls protocol ldp](#) on page 1906

[disable mpls protocol rsvp-te](#) on page 1907

[disable mpls rsvp-te bundle-message](#) on page 1907

[disable mpls rsvp-te fast-reroute](#) on page 1908

[disable mpls rsvp-te lsp](#) on page 1909

[disable mpls rsvp-te summary-refresh](#) on page 1910

[disable mpls rsvp-te](#) on page 1911

[disable mpls static lsp](#) on page 1912

[disable mpls vlan](#) on page 1912

[disable msdp](#) on page 1913

[disable msdp data-encapsulation](#) on page 1914

[disable msdp export local-sa](#) on page 1915

[disable msdp peer](#) on page 1916

[disable msdp process-sa-request](#) on page 1917

[disable msrp](#) on page 1918

[disable mvr](#) on page 1919

[disable mvrp](#) on page 1919

[disable mvrp ports](#) on page 1920

[disable neighbor-discovery refresh](#) on page 1921

[disable netlogin authentication failure vlan ports](#) on page 1922

[disable netlogin authentication service-unavailable vlan ports](#) on page 1922

[disable netlogin dot1x guest-vlan ports](#) on page 1923

[disable netlogin logout-privilege](#) on page 1924

[disable netlogin ports](#) on page 1925

[disable netlogin reauthenticate-on-refresh](#) on page 1926

[disable netlogin redirect-page](#) on page 1926

[disable netlogin session-refresh](#) on page 1927

[disable netlogin](#) on page 1928

[disable network-clock gtp ports](#) on page 1929

[disable network-clock gtp](#) on page 1929

[disable nodealias ports](#) on page 1930

[disable nodealias protocol](#) on page 1931

[disable ntp](#) on page 1932

[disable ntp authentication](#) on page 1933

[disable ntp broadcast-client](#) on page 1934
[disable ntp broadcast-server](#) on page 1935
[disable ntp vlan](#) on page 1935
[disable ntp vr](#) on page 1936
[disable ospf](#) on page 1937
[disable ospf capability opaque-lsa](#) on page 1938
[disable ospf export](#) on page 1939
[disable ospf export vr](#) on page 1940
[disable ospf mpls-next-hop](#) on page 1941
[disable ospf originate-default](#) on page 1942
[disable ospf restart-helper-lsa-check](#) on page 1942
[disable ospf use-ip-router-alert](#) on page 1943
[disable ospf vxlan-extensions](#) on page 1944
[disable ospfv3](#) on page 1945
[disable ospfv3 restart-helper-lsa-check](#) on page 1945
[disable ospfv3 export](#) on page 1946
[disable ospfv3 virtual-link restart-helper-lsa-check](#) on page 1947
[disable pim iproute sharing](#) on page 1948
[disable pim snooping](#) on page 1949
[disable pim ssm tunnel](#) on page 1950
[disable pim ssm vlan](#) on page 1951
[disable pim](#) on page 1952
[disable policy](#) on page 1952
[disable port](#) on page 1953
[disable ports mlag-id](#) on page 1954
[disable radius](#) on page 1955
[disable radius-accounting](#) on page 1956
[disable radius dynamic-authorization](#) on page 1957
[disable rip](#) on page 1957
[disable rip aggregation](#) on page 1958
[disable rip export](#) on page 1959
[disable rip originate-default](#) on page 1960
[disable rip poisonreverse](#) on page 1961
[disable rip splithorizon](#) on page 1962
[disable rip triggerupdates](#) on page 1963
[disable rip use-ip-router-alert](#) on page 1963
[disable ripng](#) on page 1964
[disable ripng export](#) on page 1965
[disable ripng originate-default](#) on page 1966
[disable ripng poisonreverse](#) on page 1967
[disable ripng splithorizon](#) on page 1968
[disable ripng triggerupdate](#) on page 1969

[disable rmon](#) on page 1970
[disable router-discovery](#) on page 1971
[disable sflow ports](#) on page 1971
[disable sflow](#) on page 1972
[disable sharing](#) on page 1973
[disable slpp guard](#) on page 1974
[disable smartredundancy](#) on page 1975
[disable snmp access vr](#) on page 1976
[disable snmp access](#) on page 1977
[disable snmp community](#) on page 1977
[disable snmp notification-log](#) on page 1978
[disable snmp trap l3vpn](#) on page 1979
[disable snmp traps](#) on page 1980
[disable snmp traps bfd](#) on page 1981
[disable snmp traps configuration](#) on page 1982
[disable snmp traps fdb mac-tracking](#) on page 1983
[disable snmp traps identity-management](#) on page 1983
[disable snmp traps l2vpn](#) on page 1984
[disable snmp traps l3vpn](#) on page 1985
[disable snmp traps lldp](#) on page 1985
[disable snmp traps lldp-med](#) on page 1986
[disable snmp traps mpls](#) on page 1987
[disable snmp traps ospf](#) on page 1988
[disable snmp traps ospfv3](#) on page 1988
[disable snmp traps port-up-down ports](#) on page 1989
[disable snmpv3](#) on page 1990
[disable snmpv3 community](#) on page 1991
[disable snmp-client](#) on page 1991
[disable ssh2](#) on page 1992
[disable stacking](#) on page 1993
[disable stacking-support](#) on page 1994
[disable stpd](#) on page 1995
[disable stpd auto-bind](#) on page 1996
[disable stpd ports](#) on page 1997
[disable stpd rapid-root-failover](#) on page 1998
[disable switch bluetooth](#) on page 1999
[disable switch locally-administered-address](#) on page 2000
[disable switch usb](#) on page 2001
[disable syslog](#) on page 2001
[disable subvlan-proxy-arp vlan](#) on page 2002
[disable tacacs](#) on page 2003
[disable tacacs-accounting](#) on page 2004

[disable tacacs-authorization](#) on page 2005
[disable tech-support collector](#) on page 2005
[disable telnet](#) on page 2006
[disable tunnel](#) on page 2007
[disable twamp reflector](#) on page 2008
[disable twamp server](#) on page 2008
[disable udp-echo-server](#) on page 2009
[disable upm profile](#) on page 2010
[disable virtual-network remote-endpoint vxlan](#) on page 2010
[disable virtual-router](#) on page 2011
[disable vlan](#) on page 2012
[disable vm autostart](#) on page 2013
[disable vm-tracking dynamic-vlan ports](#) on page 2014
[disable vm-tracking](#) on page 2015
[disable vm-tracking ports](#) on page 2016
[disable vman cep egress filtering ports](#) on page 2017
[disable vpex](#) on page 2017
[disable vpex auto-configuration](#) on page 2018
[disable vpex auto-upgrade](#) on page 2019
[disable vpls](#) on page 2020
[disable vpls fdb mac-withdrawal](#) on page 2021
[disable vpls health-check vccv](#) on page 2022
[disable vpls service](#) on page 2023
[disable vrrp group](#) on page 2024
[disable vrrp vrid](#) on page 2025
[disable watchdog](#) on page 2026
[disable web http](#) on page 2027
[disable web https](#) on page 2027
[disable cli xml-mode](#) on page 2028
[disable msrp ports](#) on page 2029
[download bootrom](#) on page 2030
[download image](#) on page 2032
[download ssl certificate](#) on page 2039
[download ssl privkey](#) on page 2041
[edit policy](#) on page 2043
[edit upm profile](#) on page 2044
[eject usb](#) on page 2045
[ELSE](#) on page 2046
[enable access-list permit to-cpu](#) on page 2047
[enable access-list refresh blackhole](#) on page 2048
[enable account](#) on page 2048
[enable avb](#) on page 2049

[enable avb ports](#) on page 2050
[enable bgp](#) on page 2051
[enable bgp advertise-inactive-route](#) on page 2052
[enable bgp aggregation](#) on page 2053
[enable bgp always-compare-med](#) on page 2054
[enable bgp community format](#) on page 2055
[enable bgp export](#) on page 2056
[enable bgp export vr](#) on page 2058
[enable bgp export \[static | direct\] l2vpn-evpn](#) on page 2059
[enable bgp fast-external-fallover](#) on page 2060
[enable bgp mpls-next-hop](#) on page 2061
[enable bgp multipath-relax](#) on page 2062
[enable bgp neighbor](#) on page 2063
[enable bgp neighbor address-family l2vpn-evpn](#) on page 2064
[enable bgp neighbor capability](#) on page 2065
[enable bgp neighbor capability address-family vpnv4](#) on page 2067
[enable bgp neighbor originate-default](#) on page 2068
[enable bgp neighbor remove-private-AS-numbers](#) on page 2070
[enable bgp neighbor soft-in-reset](#) on page 2071
[enable bgp peer-group](#) on page 2073
[enable bgp peer-group capability](#) on page 2074
[enable bgp peer-group capability](#) on page 2075
[enable bgp peer-group capability address-family vpnv4](#) on page 2077
[enable bgp peer-group next-hop-unchanged](#) on page 2078
[enable bgp peer-group originate-default](#) on page 2079
[enable bgp peer-group remove-private-AS-numbers](#) on page 2080
[enable bgp peer-group soft-in-reset](#) on page 2081
[enable bootp vlan](#) on page 2082
[enable bootprelay ipv6](#) on page 2083
[enable bootprelay](#) on page 2085
[enable cdp ports](#) on page 2086
[enable cfm segment frame-delay measurement](#) on page 2087
[enable cfm segment frame-loss measurement mep](#) on page 2088
[enable clear-flow](#) on page 2089
[enable cli history expansion](#) on page 2090
[enable cli prompting](#) on page 2091
[enable cli refresh](#) on page 2092
[enable cli scripting](#) on page 2093
[enable cli scripting output](#) on page 2094
[enable cli space-completion](#) on page 2095
[enable cli config-logging](#) on page 2095
[enable cli-config-logging expansion](#) on page 2096

[enable cli paging](#) on page 2097
[enable cpu-monitoring](#) on page 2098
[enable dhcp ports vlan](#) on page 2099
[enable dhcp vlan](#) on page 2100
[enable diffserv examination ports](#) on page 2101
[enable diffserv replacement ports](#) on page 2102
[enable dns cache](#) on page 2103
[enable dns cache dnssec](#) on page 2104
[enable dns cache analytics](#) on page 2104
[enable dos-protect simulated](#) on page 2105
[enable dos-protect](#) on page 2106
[enable dot1p examination inner-tag port](#) on page 2107
[enable dot1p examination ports](#) on page 2108
[enable dot1p replacement ports](#) on page 2109
[enable eaps](#) on page 2110
[enable edp ports](#) on page 2111
[enable elrp-client](#) on page 2112
[enable elsm ports](#) on page 2113
[enable elsm ports auto-restart](#) on page 2115
[enable erps](#) on page 2116
[enable erps block-vc-recovery](#) on page 2117
[enable erps ring-name](#) on page 2118
[enable erps topology-change](#) on page 2119
[enable esrp](#) on page 2119
[enable ethernet oam ports link-fault-management](#) on page 2120
[enable fdb static-mac-move](#) on page 2121
[enable flooding ports](#) on page 2122
[enable flow-control ports](#) on page 2123
[enable flowmon](#) on page 2125
[enable flowmon group](#) on page 2126
[enable icmp ipv6](#) on page 2127
[enable icmp redirects ipv6 fast-path](#) on page 2128
[enable icmp redirects](#) on page 2129
[enable icmp userredirects](#) on page 2129
[enable identity-management](#) on page 2130
[enable cli idle-timeout](#) on page 2131
[enable igmp](#) on page 2132
[enable igmp snooping](#) on page 2133
[enable igmp snooping vlan fast-leave](#) on page 2136
[enable igmp snooping with-proxy](#) on page 2136
[enable igmp ssm-map](#) on page 2137
[enable inline-power](#) on page 2138

[enable inline-power ports](#) on page 2140
[enable inline-power slot](#) on page 2140
[enable ip anycast](#) on page 2142
[enable ip nat](#) on page 2143
[enable ip nat rule](#) on page 2144
[enable iparp checking](#) on page 2144
[enable iparp gratuitous protect](#) on page 2145
[enable iparp refresh](#) on page 2146
[enable ipforwarding ipv6](#) on page 2147
[enable ipforwarding](#) on page 2148
[enable ipmcforwarding ipv6](#) on page 2149
[enable ipmcforwarding](#) on page 2150
[enable ipmcforwarding tunnel](#) on page 2151
[enable ip option loose-source-route](#) on page 2152
[enable ip option strict-source-route](#) on page 2153
[enable iproute bfd](#) on page 2153
[enable iproute bfd strict](#) on page 2154
[enable iproute compression](#) on page 2155
[enable iproute ipv6 compression](#) on page 2156
[enable iproute mpls-next-hop](#) on page 2157
[enable iproute protection ping](#) on page 2158
[enable iproute sharing](#) on page 2158
[enable ip-security anomaly-protection icmp](#) on page 2160
[enable ip-security anomaly-protection ip](#) on page 2160
[enable ip-security anomaly-protection l4port](#) on page 2161
[enable ip-security anomaly-protection notify](#) on page 2162
[enable ip-security anomaly-protection tcp flags](#) on page 2163
[enable ip-security anomaly-protection tcp fragment](#) on page 2163
[enable ip-security anomaly-protection](#) on page 2164
[enable ip-security arp gratuitous-protection](#) on page 2165
[enable ip-security arp learning learn-from-arp](#) on page 2166
[enable ip-security arp learning learn-from-dhcp](#) on page 2168
[enable ip-security arp validation violation-action](#) on page 2169
[enable ip-security dhcp-bindings restoration](#) on page 2171
[enable ip-security dhcp-snooping](#) on page 2172
[enable ip-security source-ip-lockdown ports](#) on page 2173
[enable iqagent](#) on page 2175
[enable irdp](#) on page 2175
[enable isis](#) on page 2176
[enable isis area adjacency-check](#) on page 2177
[enable isis area dynamic-hostname](#) on page 2178
[enable isis area export](#) on page 2179

[enable isis area export ipv6](#) on page 2180
[enable isis area originate-default](#) on page 2181
[enable isis area overload-bit](#) on page 2182
[enable isis hello-padding](#) on page 2183
[enable isis restart-helper](#) on page 2184
[enable jumbo-frame ports](#) on page 2185
[enable l2vpn](#) on page 2186
[enable l2vpn health-check vccv](#) on page 2187
[enable l2vpn service](#) on page 2188
[enable l2vpn sharing](#) on page 2189
[enable l2vpn vpls peer fdb send-mac-withdrawal](#) on page 2189
[enable learning iparp sender-mac](#) on page 2190
[enable learning port](#) on page 2191
[enable led locator](#) on page 2192
[enable license](#) on page 2193
[enable license file](#) on page 2195
[enable lldp ports](#) on page 2196
[enable log debug-mode](#) on page 2197
[enable log display](#) on page 2198
[enable log target](#) on page 2199
[enable log target upm](#) on page 2201
[enable log target xml-notification](#) on page 2202
[enable loopback-mode vlan](#) on page 2202
[enable mac-lockdown-timeout ports](#) on page 2203
[enable mac-locking ports](#) on page 2204
[enable mac-locking](#) on page 2205
[enable mirror](#) on page 2205
[enable mirror control_index](#) on page 2206
[enable mirror to port](#) on page 2207
[enable mirror to remote-ip](#) on page 2210
[enable mlag port peer id](#) on page 2214
[enable mlag port reload-delay](#) on page 2215
[enable mld](#) on page 2216
[enable mld snooping](#) on page 2217
[enable mld snooping with-proxy](#) on page 2218
[enable mld ssm-map](#) on page 2219
[enable mpls](#) on page 2219
[enable mpls bfd](#) on page 2220
[enable mpls exp examination](#) on page 2221
[enable mpls exp replacement](#) on page 2222
[enable mpls ldp bgp-routes](#) on page 2223
[enable mpls ldp loop-detection](#) on page 2224

[enable mpls ldp](#) on page 2225
[enable mpls php](#) on page 2226
[enable mpls protocol ldp](#) on page 2226
[enable mpls protocol rsvp-te](#) on page 2227
[enable mpls rsvp-te bundle-message](#) on page 2228
[enable mpls rsvp-te fast-reroute](#) on page 2229
[enable mpls rsvp-te lsp](#) on page 2230
[enable mpls rsvp-te summary-refresh](#) on page 2231
[enable mpls rsvp-te](#) on page 2232
[enable mpls static lsp](#) on page 2233
[enable mpls vlan](#) on page 2233
[enable msdp data-encapsulation](#) on page 2234
[enable msdp export local-sa](#) on page 2235
[enable msdp peer](#) on page 2237
[enable msdp process-sa-request](#) on page 2237
[enable msdp](#) on page 2239
[enable msrp ports](#) on page 2239
[enable msrp](#) on page 2240
[enable mvr](#) on page 2241
[enable mvrp](#) on page 2242
[enable mvrp ports](#) on page 2243
[enable neighbor-discovery refresh](#) on page 2243
[enable netlogin](#) on page 2244
[enable netlogin authentication failure vlan ports](#) on page 2245
[enable netlogin authentication service-unavailable vlan ports](#) on page 2246
[enable netlogin dot1x guest-vlan ports](#) on page 2247
[enable netlogin logout-privilege](#) on page 2248
[enable netlogin ports](#) on page 2249
[enable netlogin reauthentication-on-refresh](#) on page 2250
[enable netlogin redirect-page](#) on page 2250
[enable netlogin session-refresh](#) on page 2251
[enable network-clock gtp](#) on page 2252
[enable network-clock gtp ports](#) on page 2253
[enable nodealias ports](#) on page 2253
[enable nodealias protocol](#) on page 2254
[enable ntp](#) on page 2255
[enable ntp authentication](#) on page 2256
[enable ntp broadcast-client](#) on page 2257
[enable ntp broadcast-server](#) on page 2258
[enable ntp vlan](#) on page 2259
[enable ntp vr](#) on page 2260
[enable ospf](#) on page 2260

[enable ospf capability opaque-lsa](#) on page 2261
[enable ospf export](#) on page 2262
[enable ospf export vr](#) on page 2264
[enable ospf mpls-next-hop](#) on page 2265
[enable ospf originate-default](#) on page 2266
[enable ospf restart-helper-lsa-check](#) on page 2267
[enable ospf use-ip-router-alert](#) on page 2268
[enable ospfv3](#) on page 2268
[enable ospfv3 export](#) on page 2269
[enable ospfv3 restart-helper-lsa-check](#) on page 2271
[enable ospfv3 virtual-link restart-helper-lsa-check](#) on page 2272
[enable pim](#) on page 2273
[enable pim iproute sharing](#) on page 2274
[enable pim snooping](#) on page 2275
[enable pim ssm tunnel](#) on page 2275
[enable pim ssm vlan](#) on page 2276
[enable policy](#) on page 2277
[enable port](#) on page 2278
[enable ports mlag-id](#) on page 2279
[enable radius](#) on page 2280
[enable radius-accounting](#) on page 2281
[enable radius dynamic-authorization](#) on page 2282
[enable rip](#) on page 2283
[enable rip aggregation](#) on page 2284
[enable rip export](#) on page 2285
[enable rip originate-default cost](#) on page 2286
[enable rip poisonreverse](#) on page 2287
[enable rip splithorizon](#) on page 2288
[enable rip triggerupdates](#) on page 2289
[enable rip use-ip-router-alert](#) on page 2290
[enable ripng](#) on page 2291
[enable ripng export](#) on page 2292
[enable ripng originate-default](#) on page 2293
[enable ripng poisonreverse](#) on page 2294
[enable ripng splithorizon](#) on page 2295
[enable ripng triggerupdates](#) on page 2296
[enable rmon](#) on page 2297
[enable router-discovery](#) on page 2299
[enable sflow](#) on page 2300
[enable sflow ports](#) on page 2301
[enable sharing grouping](#) on page 2302
[enable slpp guard](#) on page 2305

[enable smartredundancy](#) on page 2306
[enable snmp access](#) on page 2307
[enable snmp access vr](#) on page 2308
[enable snmp community](#) on page 2309
[enable snmp notification-log](#) on page 2310
[enable snmp trap l3vpn](#) on page 2311
[enable snmp traps](#) on page 2312
[enable snmp traps configuration](#) on page 2313
[enable snmp traps bfd](#) on page 2313
[enable snmp traps fdb mac-tracking](#) on page 2314
[enable snmp traps identity-management](#) on page 2315
[enable snmp traps l2vpn](#) on page 2316
[enable snmp traps l3vpn](#) on page 2316
[enable snmp traps lldp](#) on page 2317
[enable snmp traps lldp-med](#) on page 2318
[enable snmp traps mpls](#) on page 2319
[enable snmp traps ospf](#) on page 2320
[enable snmp traps ospfv3](#) on page 2320
[enable snmp traps port-up-down ports](#) on page 2321
[enable snmpv3](#) on page 2322
[enable snmpv3 community](#) on page 2323
[enable snmp-client](#) on page 2324
[enable ssh2](#) on page 2324
[enable stacking](#) on page 2326
[enable stacking-support](#) on page 2328
[enable stpd](#) on page 2329
[enable stpd auto-bind](#) on page 2329
[enable stpd ports](#) on page 2332
[enable stpd rapid-root-failover](#) on page 2333
[enable subvlan-proxy-arp vlan](#) on page 2334
[enable switch bluetooth](#) on page 2335
[enable switch locally-administered-address](#) on page 2336
[enable switch usb](#) on page 2337
[enable syslog](#) on page 2338
[enable tacacs](#) on page 2339
[enable tacacs-accounting](#) on page 2339
[enable tacacs-authorization](#) on page 2340
[enable tech-support collector](#) on page 2341
[enable telnet](#) on page 2342
[enable tunnel](#) on page 2343
[enable twamp reflector](#) on page 2344
[enable twamp server](#) on page 2344

[enable udp-echo-server](#) on page 2345
[enable upm profile](#) on page 2346
[enable virtual-network remote-endpoint vxlan](#) on page 2346
[enable virtual-router](#) on page 2347
[enable vlan](#) on page 2348
[enable vman cep egress filtering ports](#) on page 2349
[enable vm autostart](#) on page 2350
[enable vm-tracking](#) on page 2351
[enable vm-tracking dynamic-vlan ports](#) on page 2352
[enable vm-tracking ports](#) on page 2352
[enable vpex](#) on page 2353
[enable vpex auto-configuration](#) on page 2354
[enable vpex auto-upgrade](#) on page 2355
[enable vpls](#) on page 2356
[enable vpls fdb mac-withdrawal](#) on page 2357
[enable vpls health-check vccv](#) on page 2358
[enable vpls service](#) on page 2359
[enable vrrp group](#) on page 2360
[enable vrrp vrid](#) on page 2361
[enable watchdog](#) on page 2362
[enable web http](#) on page 2363
[enable web https](#) on page 2363
[enable cli xml-mode](#) on page 2364
[enable/disable bfd vlan](#) on page 2365
[enable/disable xml-notification](#) on page 2366
[ENDIF](#) on page 2367
[ENDWHILE](#) on page 2368
[exit](#) on page 2369
[history](#) on page 2370
[IF ... THEN](#) on page 2371
[install bootrom](#) on page 2372
[install firmware](#) on page 2374
[install image](#) on page 2376
[install image inactive](#) on page 2379
[install license file](#) on page 2380
[load script](#) on page 2381
[load var key](#) on page 2383
[logout](#) on page 2384
[ls](#) on page 2384
[mkdir](#) on page 2387
[mrinfo](#) on page 2388
[mtrace](#) on page 2389

[mv](#) on page 2392
[nslookup](#) on page 2394
[open vm console](#) on page 2396
[ping](#) on page 2397
[ping mac port](#) on page 2399
[ping mpls lsp](#) on page 2400
[pwd](#) on page 2402
[quit](#) on page 2403
[reboot](#) on page 2404
[refresh access-list network-zone](#) on page 2406
[refresh identity-management role](#) on page 2407
[refresh igmp ssm-map](#) on page 2408
[refresh mld ssm-map](#) on page 2409
[refresh policy](#) on page 2410
[reset inline-power ports](#) on page 2411
[restart ports](#) on page 2412
[restart process](#) on page 2413
[restart process mpls](#) on page 2415
[restart vm](#) on page 2415
[resume vm](#) on page 2416
[return](#) on page 2417
[rm](#) on page 2418
[rmdir](#) on page 2420
[rtlookup rpf](#) on page 2420
[rtlookup](#) on page 2421
[run diagnostics](#) on page 2422
[run elrp](#) on page 2425
[run failover](#) on page 2426
[run script](#) on page 2427
[run tech-support report](#) on page 2429
[run update](#) on page 2430
[run upm profile](#) on page 2431
[run vm-tracking repository](#) on page 2432
[save configuration](#) on page 2433
[save configuration as-script](#) on page 2436
[save configuration automatic](#) on page 2437
[save debug tracefiles](#) on page 2439
[save var key](#) on page 2440
[save vm image](#) on page 2441
[save vm state](#) on page 2441
[scp2](#) on page 2443
[set var](#) on page 2446

[show access-list](#) on page 2447
[show access-list configuration](#) on page 2448
[show access-list counter](#) on page 2450
[show access-list counters process](#) on page 2451
[show access-list dynamic rule](#) on page 2452
[show access-list dynamic counter](#) on page 2454
[show access-list dynamic](#) on page 2455
[show access-list interface](#) on page 2456
[show access-list meter](#) on page 2458
[show access-list network-zone](#) on page 2459
[show access-list usage acl-mask port](#) on page 2460
[show access-list usage acl-range port](#) on page 2461
[show access-list usage acl-rule port](#) on page 2462
[show access-list usage acl-slice port](#) on page 2463
[show access-list width](#) on page 2465
[show accounts](#) on page 2466
[show accounts password-policy](#) on page 2467
[show auto-peering](#) on page 2469
[show auto-peering one-config](#) on page 2470
[show auto-provision](#) on page 2472
[show automation edge database](#) on page 2474
[show automation edge devices](#) on page 2475
[show avb](#) on page 2476
[show bandwidth pool](#) on page 2477
[show banner](#) on page 2478
[show banner netlogin](#) on page 2479
[show bfd](#) on page 2480
[show bfd counters](#) on page 2481
[show bfd session client](#) on page 2482
[show bfd session counters missed-hellos](#) on page 2483
[show bfd session counters vr all](#) on page 2484
[show bfd session detail vr all](#) on page 2485
[show bfd session vr all](#) on page 2487
[show bfd vlan counters](#) on page 2488
[show bfd vlan](#) on page 2489
[show bgp](#) on page 2490
[show bgp evpn](#) on page 2494
[show bpg evpn evi](#) on page 2495
[show bgp evpn ipv4](#) on page 2496
[show bgp evpn ipv6](#) on page 2498
[show bgp evpn l3vni](#) on page 2499
[show bgp evpn mac](#) on page 2500

[show bgp memory](#) on page 2502
[show bgp neighbor \[flap-statistics | suppressed-routes\]](#) on page 2504
[show bgp neighbor received orf](#) on page 2507
[show bgp neighbor](#) on page 2508
[show bgp peer-group](#) on page 2515
[show bgp routes summary](#) on page 2518
[show bgp routes](#) on page 2520
[show bootprelay](#) on page 2526
[show bootprelay configuration](#) on page 2527
[show bootprelay configuration ipv4](#) on page 2529
[show bootprelay configuration ipv6](#) on page 2531
[show bootprelay dhcp-agent information circuit-id port-information](#) on page 2532
[show bootprelay dhcp-agent information circuit-id vlan-information](#) on page 2533
[show bootprelay ipv6](#) on page 2534
[show bootprelay ipv6 prefix-delegation snooping](#) on page 2535
[show cdp](#) on page 2536
[show cdp counters](#) on page 2537
[show cdp neighbor](#) on page 2538
[show cdp ports](#) on page 2538
[show cfm detail](#) on page 2540
[show cfm groups](#) on page 2542
[show cfm segment frame-delay statistics](#) on page 2545
[show cfm segment frame-delay](#) on page 2546
[show cfm segment frame-delay/frame-loss mep id](#) on page 2547
[show cfm segment frame-loss statistics](#) on page 2550
[show cfm segment frame-loss](#) on page 2551
[show cfm segment mep](#) on page 2552
[show cfm segment](#) on page 2554
[show cfm session counters missed-hellos](#) on page 2556
[show cfm](#) on page 2557
[show checkpoint-data](#) on page 2560
[show clear-flow](#) on page 2562
[show clear-flow acl-modified](#) on page 2563
[show clear-flow rule](#) on page 2564
[show clear-flow rule-all](#) on page 2565
[show clear-flow rule-triggered](#) on page 2567
[show cli journal](#) on page 2567
[show configuration](#) on page 2569
[show configuration difference](#) on page 2572
[show configuration “xmlc”](#) on page 2574
[show cos-index](#) on page 2575
[show counters vr](#) on page 2576

[show cpu-monitoring](#) on page 2577
[show database](#) on page 2580
[show debug](#) on page 2581
[show dhcp-client state](#) on page 2582
[show dhcp-server](#) on page 2583
[show diagnostics](#) on page 2584
[show diffserv examination](#) on page 2586
[show diffserv replacement](#) on page 2587
[show dns-client](#) on page 2588
[show dns cache analytics configuration](#) on page 2588
[show dns cache analytics protected-client](#) on page 2589
[show dns cache configuration](#) on page 2590
[show dns cache analytics statistics](#) on page 2591
[show dns cache](#) on page 2594
[show dns cache name-server](#) on page 2596
[show dos-protect](#) on page 2596
[show dot1p](#) on page 2598
[show dwdm channel-map](#) on page 2599
[show eaps](#) on page 2600
[show eaps cfm groups](#) on page 2605
[show eaps counters shared-port](#) on page 2606
[show eaps counters](#) on page 2610
[show eaps shared-port](#) on page 2615
[show eaps shared-port neighbor-info](#) on page 2619
[show edp](#) on page 2620
[show elrp](#) on page 2623
[show elrp disabled-ports](#) on page 2625
[show elrp dynamic-vlans](#) on page 2626
[show elsm ports](#) on page 2628
[show elsm](#) on page 2632
[show erps](#) on page 2635
[show erps ring-name](#) on page 2636
[show erps statistics](#) on page 2637
[show esrp](#) on page 2638
[show esrp aware](#) on page 2640
[show esrp counters](#) on page 2642
[show ethernet oam](#) on page 2643
[show fabric attach agent](#) on page 2645
[show fabric attach assignments](#) on page 2646
[show fabric attach elements](#) on page 2647
[show fabric attach port](#) on page 2648
[show fabric attach ports authentication](#) on page 2648

[show fabric attach statistics](#) on page 2650
[show fabric attach zero-touch-client](#) on page 2650
[show failsafe-account](#) on page 2651
[show fans](#) on page 2652
[show fdb](#) on page 2654
[show fdb mac-tracking configuration](#) on page 2658
[show fdb mac-tracking statistics](#) on page 2658
[show fdb static-mac-move configuration](#) on page 2660
[show fdb stats](#) on page 2660
[show flow-redirect](#) on page 2662
[show flowmon](#) on page 2664
[show flowmon collector](#) on page 2665
[show flowmon group](#) on page 2666
[show flowmon group statistics](#) on page 2667
[show flowmon group template](#) on page 2668
[show flowmon key](#) on page 2669
[show forwarding configuration](#) on page 2670
[show forwarding hardware-utilization](#) on page 2672
[show heartbeat process](#) on page 2674
[show identity-management blacklist](#) on page 2675
[show identity-management entries](#) on page 2676
[show identity-management greylist](#) on page 2680
[show identity-management list-precedence](#) on page 2681
[show identity-management role](#) on page 2682
[show identity-management statistics](#) on page 2684
[show identity-management whitelist](#) on page 2685
[show identity-management](#) on page 2686
[show igmp](#) on page 2687
[show igmp counters](#) on page 2688
[show igmp group](#) on page 2689
[show igmp snooping cache](#) on page 2690
[show igmp snooping vlan filter](#) on page 2691
[show igmp snooping vlan static](#) on page 2692
[show igmp snooping vlan](#) on page 2693
[show igmp snooping](#) on page 2694
[show igmp ssm-map](#) on page 2695
[show inline-power configuration ports](#) on page 2697
[show inline-power fast ports](#) on page 2698
[show inline-power info ports](#) on page 2699
[show inline-power slot](#) on page 2703
[show inline-power stats ports](#) on page 2705
[show inline-power stats slot](#) on page 2707

[show inline-power stats](#) on page 2708
[show inline-power](#) on page 2709
[show ip nat](#) on page 2710
[show ip nat rule](#) on page 2711
[show ip nat rule statistics](#) on page 2714
[show ip nat vlan](#) on page 2715
[show ip nat vlan counters](#) on page 2716
[show iparp](#) on page 2717
[show iparp proxy](#) on page 2719
[show iparp security](#) on page 2720
[show iparp stats](#) on page 2722
[show ipconfig](#) on page 2724
[show ipconfig ipv6](#) on page 2725
[show ipmroute](#) on page 2727
[show iproute](#) on page 2728
[show iproute bfd](#) on page 2731
[show iproute ipv6 origin](#) on page 2732
[show iproute ipv6](#) on page 2733
[show iproute mpls origin](#) on page 2735
[show iproute mpls](#) on page 2737
[show iproute multicast](#) on page 2738
[show iproute origin](#) on page 2739
[show iproute protection ping](#) on page 2741
[show iproute reserved-entries statistics](#) on page 2742
[show iproute reserved-entries](#) on page 2744
[show ip-security anomaly-protection notify cache ports](#) on page 2745
[show ip-security arp gratuitous-protection](#) on page 2746
[show ip-security arp learning](#) on page 2746
[show ip-security arp validation](#) on page 2748
[show ip-security arp validation violations](#) on page 2749
[show ip-security dhcp-snooping entries](#) on page 2750
[show ip-security dhcp-snooping information circuit-id port-information](#) on page 2751
[show ip-security dhcp-snooping information-option](#) on page 2752
[show ip-security dhcp-snooping information-option circuit-id vlan-information](#) on page 2753
[show ip-security dhcp-snooping information remote-id](#) on page 2753
[show ip-security dhcp-snooping](#) on page 2755
[show ip-security dhcp-snooping violations](#) on page 2756
[show ip-security source-ip-lockdown](#) on page 2757
[show ipstats ipv6](#) on page 2758
[show ipstats](#) on page 2759
[show ipv6 dad](#) on page 2760
[show iqagent](#) on page 2762

[show iqagent discovery detail](#) on page 2764
[show isis](#) on page 2765
[show isis area summary-addresses](#) on page 2766
[show isis area](#) on page 2767
[show isis counters](#) on page 2768
[show isis lsdb](#) on page 2769
[show isis neighbors](#) on page 2770
[show isis topology](#) on page 2771
[show isis vlan](#) on page 2772
[show keychain](#) on page 2773
[show l2pt](#) on page 2775
[show l2pt profile](#) on page 2775
[show L2stats](#) on page 2776
[show l2vpn](#) on page 2777
[show lacp](#) on page 2783
[show lacp counters](#) on page 2784
[show lacp lag](#) on page 2786
[show lacp member-port](#) on page 2789
[show ldap domain](#) on page 2792
[show ldap statistics](#) on page 2795
[show licenses](#) on page 2797
[show lldp](#) on page 2799
[show lldp dcbx](#) on page 2801
[show lldp neighbors](#) on page 2807
[show lldp statistics](#) on page 2810
[show log](#) on page 2811
[show log components](#) on page 2815
[show log configuration filter](#) on page 2819
[show log configuration target](#) on page 2820
[show log configuration target upm](#) on page 2824
[show log configuration target xml-notification](#) on page 2826
[show log configuration](#) on page 2828
[show log counters](#) on page 2831
[show log events](#) on page 2834
[show mac-lockdown-timeout fdb ports](#) on page 2835
[show mac-lockdown-timeout ports](#) on page 2837
[show mac-locking stations](#) on page 2838
[show mac-locking](#) on page 2839
[show macsec](#) on page 2840
[show macsec connectivity-association](#) on page 2842
[show macsec encryption-engine monitor](#) on page 2843
[show macsec ports](#) on page 2844

[show macsec ports configuration](#) on page 2847
[show macsec ports detail](#) on page 2848
[show macsec ports usage](#) on page 2851
[show macsec usage](#) on page 2853
[show management](#) on page 2854
[show mcast cache](#) on page 2857
[show mcast ipv6 cache](#) on page 2859
[show memory](#) on page 2861
[show memory process](#) on page 2864
[show meter](#) on page 2866
[show meter out-of-profile](#) on page 2868
[show mirror](#) on page 2869
[show mlag peer](#) on page 2872
[show mlag ports](#) on page 2874
[show mld](#) on page 2876
[show mld counters](#) on page 2877
[show mld group](#) on page 2878
[show mld snooping](#) on page 2879
[show mld snooping vlan filter](#) on page 2881
[show mld snooping vlan static](#) on page 2881
[show mld ssm-map](#) on page 2882
[show mpls](#) on page 2884
[show mpls bfd](#) on page 2885
[show mpls exp examination](#) on page 2886
[show mpls exp replacement](#) on page 2887
[show mpls interface](#) on page 2888
[show mpls label](#) on page 2889
[show mpls label l3vpn](#) on page 2892
[show mpls label usage](#) on page 2894
[show mpls ldp](#) on page 2896
[show mpls ldp interface](#) on page 2898
[show mpls ldp label](#) on page 2900
[show mpls ldp label advertised](#) on page 2901
[show mpls ldp label l2vpn retained](#) on page 2902
[show mpls ldp label l2vpn](#) on page 2904
[show mpls ldp label lsp retained](#) on page 2905
[show mpls ldp label retained](#) on page 2906
[show mpls ldp lsp](#) on page 2908
[show mpls ldp peer](#) on page 2909
[show mpls rsvp-te bandwidth](#) on page 2912
[show mpls rsvp-te interface](#) on page 2914
[show mpls rsvp-te lsp](#) on page 2916

[show mpls rsvp-te lsp \[egress | transit\]](#) on page 2919
[show mpls rsvp-te lsp ingress](#) on page 2920
[show mpls rsvp-te neighbor](#) on page 2923
[show mpls rsvp-te path](#) on page 2924
[show mpls rsvp-te profile](#) on page 2925
[show mpls rsvp-te profile fast-reroute](#) on page 2927
[show mpls rsvp-te](#) on page 2928
[show mpls static lsp](#) on page 2929
[show mpls statistics l2vpn](#) on page 2930
[show mrp ports](#) on page 2932
[show msdp memory](#) on page 2933
[show msdp mesh-group](#) on page 2934
[show msdp peer](#) on page 2935
[show msdp sa-cache](#) on page 2937
[show msdp](#) on page 2939
[show msrp](#) on page 2939
[show msrp listeners](#) on page 2940
[show msrp ports](#) on page 2942
[show msrp ports bandwidth](#) on page 2944
[show msrp ports counters](#) on page 2945
[show msrp streams](#) on page 2946
[show msrp talkers](#) on page 2948
[show mvr](#) on page 2950
[show mvr cache](#) on page 2951
[show mvrp](#) on page 2952
[show mvrp ports counters](#) on page 2952
[show mvrp tag](#) on page 2954
[show netlogin](#) on page 2955
[show neighbor-discovery cache ipv6](#) on page 2960
[show netlogin authentication failure vlan](#) on page 2961
[show netlogin authentication service-unavailable vlan](#) on page 2962
[show netlogin banner](#) on page 2963
[show netlogin guest-vlan](#) on page 2964
[show netlogin local-users](#) on page 2965
[show netlogin mac-list](#) on page 2966
[show netlogin session](#) on page 2967
[show netlogin timeout](#) on page 2969
[show netlogin trap](#) on page 2970
[show network-clock gtp](#) on page 2971
[show network-clock gtp ports](#) on page 2972
[show node](#) on page 2976
[show nodealias](#) on page 2978

[show nodealias ip address](#) on page 2979
[show nodealias mac](#) on page 2980
[show nodealias ports](#) on page 2981
[show nodealias protocol](#) on page 2982
[show ntp](#) on page 2983
[show ntp association statistics](#) on page 2983
[show ntp association](#) on page 2985
[show ntp key](#) on page 2986
[show ntp restrict-list](#) on page 2987
[show ntp server](#) on page 2988
[show ntp sys-info](#) on page 2989
[show ntp vlan](#) on page 2990
[show ntp vr](#) on page 2991
[show odometers](#) on page 2991
[show ospf](#) on page 2993
[show ospf area](#) on page 2994
[show ospf ase-summary](#) on page 2995
[show ospf inter-vr-export](#) on page 2995
[show ospf inter-vr-export detail](#) on page 2996
[show ospf interfaces](#) on page 2997
[show ospf interfaces detail](#) on page 2999
[show ospf lsdb](#) on page 3000
[show ospf memory](#) on page 3001
[show ospf neighbor](#) on page 3002
[show ospf virtual-link](#) on page 3003
[show ospfv3](#) on page 3004
[show ospfv3 area](#) on page 3005
[show ospfv3 interfaces](#) on page 3007
[show ospfv3 lsdb stats](#) on page 3010
[show ospfv3 lsdb](#) on page 3011
[show ospfv3 neighbor](#) on page 3013
[show ospfv3 virtual-link](#) on page 3014
[show pim anycast-rp](#) on page 3016
[show pim cache](#) on page 3018
[show pim](#) on page 3020
[show pim snooping](#) on page 3025
[show pim tunnel](#) on page 3026
[show policy](#) on page 3027
[show policy access-list](#) on page 3028
[show policy access-list action-set](#) on page 3032
[show policy allowed-type](#) on page 3032
[show policy app-signature](#) on page 3033

[show policy app-signature group](#) on page 3034
[show policy autoclear](#) on page 3036
[show policy capability](#) on page 3037
[show policy captive-portal](#) on page 3039
[show policy convergence-endpoint](#) on page 3042
[show policy convergence-endpoint connections](#) on page 3043
[show policy convergence-endpoint ports](#) on page 3043
[show policy dynamic](#) on page 3044
[show policy invalid](#) on page 3045
[show policy maptable](#) on page 3046
[show policy profile](#) on page 3046
[show policy resource-profile](#) on page 3048
[show policy rule](#) on page 3050
[show policy rule port-hit](#) on page 3054
[show policy slices](#) on page 3056
[show policy state](#) on page 3056
[show policy syslog](#) on page 3057
[show policy vlanauthorization](#) on page 3058
[show ports](#) on page 3059
[show ports advertised](#) on page 3061
[show ports anomaly](#) on page 3062
[show ports buffer](#) on page 3064
[show ports collisions](#) on page 3065
[show ports configuration](#) on page 3066
[show ports congestion](#) on page 3069
[show ports eee](#) on page 3071
[show ports flow-control](#) on page 3072
[show port forward-error-correction](#) on page 3074
[show ports group](#) on page 3075
[show port information](#) on page 3075
[show ports link-flap-detection](#) on page 3085
[show ports link-scan](#) on page 3087
[show ports macsec-engines](#) on page 3087
[show ports packet](#) on page 3089
[show ports partition-template](#) on page 3090
[show ports protocol filter](#) on page 3091
[show ports qosmonitor](#) on page 3093
[show ports qosmonitor {congestion}](#) on page 3095
[show ports rate-limit flood](#) on page 3097
[show ports redundant](#) on page 3100
[show ports rxerrors](#) on page 3101
[show ports sharing](#) on page 3103

[show ports stack-ports congestion](#) on page 3105
[show ports stack-ports qosmonitor](#) on page 3106
[show ports stack-ports qosmonitor congestion](#) on page 3107
[show ports statistics](#) on page 3109
[show ports transceiver information detail](#) on page 3112
[show ports transceiver information](#) on page 3115
[show ports txerrors](#) on page 3116
[show ports utilization](#) on page 3119
[show ports vlan statistics](#) on page 3121
[show ports wred](#) on page 3122
[show power](#) on page 3125
[show power \(Stack Nodes Only\)](#) on page 3128
[show private-vlan](#) on page 3129
[show private-vlan name](#) on page 3131
[show process](#) on page 3132
[show process group](#) on page 3137
[show protocol](#) on page 3139
[show qosprofile](#) on page 3140
[show qosscheduler](#) on page 3142
[show radius](#) on page 3143
[show radius-accounting](#) on page 3146
[show radius dynamic-authorization](#) on page 3147
[show rip](#) on page 3148
[show rip interface vlan](#) on page 3149
[show rip interface](#) on page 3150
[show rip memory](#) on page 3152
[show rip routes](#) on page 3153
[show ripng](#) on page 3154
[show ripng interface](#) on page 3155
[show ripng routes](#) on page 3156
[show rmon memory](#) on page 3158
[show router-discovery](#) on page 3160
[show rtep l2pt](#) on page 3161
[show script output autoexec](#) on page 3163
[show script output default](#) on page 3164
[show security](#) on page 3165
[show session](#) on page 3170
[show sflow configuration](#) on page 3172
[show sflow hardware-utilization](#) on page 3173
[show sflow statistics](#) on page 3175
[show sharing](#) on page 3176
[show sharing distribution port-based](#) on page 3177

[show sharing health-check](#) on page 3178
[show sharing port-based keys](#) on page 3179
[show slot](#) on page 3180
[show slpp guard](#) on page 3184
[show snmp](#) on page 3187
[show snmp notification-log entry](#) on page 3187
[show snmp notification-log name](#) on page 3189
[show snmp notification-log](#) on page 3190
[show snmp traps bfd](#) on page 3191
[show snmp traps configuration](#) on page 3192
[show snmp vr_name](#) on page 3193
[show snmpv3 access](#) on page 3194
[show snmpv3 community](#) on page 3196
[show snmpv3 context](#) on page 3197
[show snmpv3 counters](#) on page 3198
[show snmpv3 engine-info](#) on page 3199
[show snmpv3 extreme-target-addr-ext](#) on page 3200
[show snmpv3 filter](#) on page 3201
[show snmpv3 filter-profile](#) on page 3202
[show snmpv3 group](#) on page 3203
[show snmpv3 mib-view](#) on page 3206
[show snmpv3 notify](#) on page 3207
[show snmpv3 target-addr](#) on page 3209
[show snmpv3 target-params](#) on page 3210
[show snmpv3 user](#) on page 3211
[show sntp-client](#) on page 3213
[show ssh2](#) on page 3214
[show ssh2 ciphers macs](#) on page 3216
[show ssh2 private-key](#) on page 3216
[show sshd2 user-key](#) on page 3217
[show ssl](#) on page 3218
[show ssl csr](#) on page 3219
[show stack-ports debounce](#) on page 3220
[show stacking](#) on page 3221
[show stacking configuration](#) on page 3224
[show stacking detail](#) on page 3225
[show stacking stack-ports](#) on page 3228
[show stacking-support](#) on page 3230
[show stpd ports blocked-ports](#) on page 3232
[show stpd ports counters](#) on page 3233
[show stpd ports non-forwarding-reason](#) on page 3235
[show stpd](#) on page 3236

[show stpd ports](#) on page 3240
[show switch](#) on page 3243
[show switch bluetooth](#) on page 3246
[show switch management](#) on page 3248
[show switch mounts](#) on page 3249
[show switch usb](#) on page 3251
[show system](#) on page 3251
[show tacacs](#) on page 3254
[show tacacs-accounting](#) on page 3256
[show time](#) on page 3257
[show tech-support](#) on page 3258
[show tech-support collector](#) on page 3260
[show temperature](#) on page 3262
[show tunnel](#) on page 3263
[show twamp endpoint](#) on page 3266
[show twamp reflector](#) on page 3267
[Object Missing](#) on page 3268
[show upm event](#) on page 3268
[show upm history](#) on page 3269
[show upm history exec-id](#) on page 3270
[show upm profile](#) on page 3271
[show upm timers](#) on page 3272
[show var](#) on page 3273
[show version](#) on page 3274
[show virtual-network](#) on page 3276
[show virtual-network remote-endpoint vxlan](#) on page 3278
[show virtual-network statistics](#) on page 3280
[show virtual-router](#) on page 3281
[show vlan](#) on page 3286
[show vlan description](#) on page 3291
[show vlan dhcp-config](#) on page 3292
[show vlan dhcp-address-allocation](#) on page 3293
[show vlan dynamic-vlan](#) on page 3294
[show vlan eaps](#) on page 3295
[show vlan fabric attach assignments](#) on page 3296
[show vlan l2pt](#) on page 3297
[show vlan security](#) on page 3298
[show vlan statistics](#) on page 3299
[show vlan stpd](#) on page 3300
[show vm](#) on page 3302
[show vm guest interfaces](#) on page 3304
[show vm virtual-interface](#) on page 3305

[show vman](#) on page 3306
[show vman eaps](#) on page 3308
[show vman ethertype](#) on page 3309
[show vm-tracking](#) on page 3310
[show vm-tracking local-vm](#) on page 3311
[show vm-tracking network-vm](#) on page 3312
[show vm-tracking nms](#) on page 3313
[show vm-tracking port](#) on page 3314
[show vm-tracking repository](#) on page 3315
[show vm-tracking vpp](#) on page 3316
[show vpex](#) on page 3317
[show vpex auto-configuration](#) on page 3319
[show vpex bpe](#) on page 3320
[show vpex bpe cpu-utilization](#) on page 3320
[show vpex bpe environment](#) on page 3321
[show vpex bpe statistics](#) on page 3323
[show vpex bpe version detail](#) on page 3326
[show vpex ports](#) on page 3327
[show vpex ports ecp statistics](#) on page 3328
[show vpex ports statistics](#) on page 3329
[show vpex stacking](#) on page 3332
[show vpex topology](#) on page 3333
[show vpls](#) on page 3335
[show vpls peer l2pt](#) on page 3343
[show vrrp](#) on page 3344
[show vrrp group](#) on page 3347
[show vrrp vlan](#) on page 3348
[show wredprofile](#) on page 3350
[show xml-notification configuration](#) on page 3351
[show xml-notification statistics](#) on page 3352
[ssh2](#) on page 3353
[start orchestration mlag](#) on page 3356
[start process](#) on page 3357
[start vm](#) on page 3359
[stop orchestration](#) on page 3360
[stop vm](#) on page 3361
[suspend vm](#) on page 3362
[synchronize](#) on page 3362
[synchronize stacking](#) on page 3365
[telnet slot](#) on page 3366
[telnet](#) on page 3367
[terminate process](#) on page 3369

[terminate vpex ztp](#) on page 3370

[tftp](#) on page 3371

[tftp get](#) on page 3375

[tftp put](#) on page 3377

[top](#) on page 3379

[traceroute](#) on page 3380

[traceroute mac port](#) on page 3381

[traceroute mpls lsp](#) on page 3383

[unalias](#) on page 3385

[unconfigure access-list](#) on page 3386

[unconfigure avb](#) on page 3387

[unconfigure banner](#) on page 3388

[unconfigure bfd vlan](#) on page 3389

[unconfigure bootprelay dhcp-agent information check](#) on page 3390

[unconfigure bootprelay dhcp-agent information circuit-id port-information](#) on page 3391

[unconfigure bootprelay dhcp-agent information circuit-id vlan-information](#) on page 3391

[unconfigure bootprelay dhcp-agent information option](#) on page 3392

[unconfigure bootprelay dhcp-agent information policy](#) on page 3393

[unconfigure bootprelay dhcp-agent information remote-id](#) on page 3394

[unconfigure bootprelay dhcp-agent source-vlan](#) on page 3394

[unconfigure bootprelay include-secondary](#) on page 3395

[unconfigure cfm domain association end-point transmit-interval](#) on page 3396

[unconfigure cos-index](#) on page 3397

[unconfigure diffserv examination](#) on page 3398

[unconfigure diffserv replacement](#) on page 3399

[unconfigure eaps port](#) on page 3399

[unconfigure eaps shared-port link-id](#) on page 3401

[unconfigure eaps shared-port mode](#) on page 3401

[unconfigure elrp-client](#) on page 3402

[unconfigure elrp-client disable ports](#) on page 3403

[unconfigure erps cfm](#) on page 3404

[unconfigure erps neighbor-port](#) on page 3405

[unconfigure erps notify-topology-change](#) on page 3405

[unconfigure erps protection-port](#) on page 3406

[unconfigure erps ring-ports west](#) on page 3407

[unconfigure icmp](#) on page 3408

[unconfigure igmp](#) on page 3409

[unconfigure identity-management list-precedence](#) on page 3409

[unconfigure identity-management](#) on page 3410

[unconfigure igmp snooping vlan ports set join-limit](#) on page 3411

[unconfigure igmp ssm-map](#) on page 3412

[unconfigure inline-power classification](#) on page 3413

[unconfigure inline-power detection ports](#) on page 3414

[unconfigure inline-power disconnect-precedence](#) on page 3415

[unconfigure inline-power operator-limit ports](#) on page 3416

[unconfigure inline-power priority ports](#) on page 3417

[unconfigure inline-power usage-threshold](#) on page 3418

[unconfigure iparp](#) on page 3419

[unconfigure ip-fix](#) on page 3419

[unconfigure ip-fix flow-key](#) on page 3420

[unconfigure ip-fix ip-address](#) on page 3421

[unconfigure ip-fix ports](#) on page 3422

[unconfigure ip-fix ports flow-key mask](#) on page 3423

[unconfigure ip-fix source ip-address](#) on page 3423

[unconfigure iproute priority](#) on page 3424

[unconfigure iproute ipv6 priority](#) on page 3427

[unconfigure ip-security dhcp-bindings storage filename](#) on page 3428

[unconfigure ip-security dhcp-snooping information check](#) on page 3429

[unconfigure ip-security dhcp-snooping information circuit-id port-information ports](#) on page 3430

[unconfigure ip-security dhcp-snooping information circuit-id vlan-information](#) on page 3430

[unconfigure ip-security dhcp-snooping information option](#) on page 3431

[unconfigure ip-security dhcp-snooping information policy](#) on page 3432

[unconfigure ip-security dhcp-snooping information remote-id](#) on page 3432

[Object Missing](#) on page 3433

[unconfigure isis area](#) on page 3433

[unconfigure isis vlan](#) on page 3434

[unconfigure l2vpn dot1q ethertype](#) on page 3435

[unconfigure l2vpn vpls redundancy](#) on page 3436

[unconfigure ldap domains](#) on page 3437

[unconfigure lldp](#) on page 3437

[unconfigure log filter](#) on page 3438

[unconfigure log target format](#) on page 3439

[unconfigure meter](#) on page 3441

[unconfigure mlag peer interval](#) on page 3442

[unconfigure mlag peer ipaddress](#) on page 3442

[unconfigure mld](#) on page 3443

[unconfigure mld ssm-map](#) on page 3444

[unconfigure mpls exp examination](#) on page 3445

[unconfigure mpls exp replacement](#) on page 3445

[unconfigure mpls vlan](#) on page 3446

[unconfigure mpls](#) on page 3447

[unconfigure mrp ports timers](#) on page 3448

[unconfigure msdp sa-cache-server](#) on page 3449

[unconfigure msrp](#) on page 3450
[unconfigure mstp region](#) on page 3451
[unconfigure mvrp stpd](#) on page 3452
[unconfigure mvrp tag](#) on page 3453
[unconfigure mvrp](#) on page 3453
[unconfigure neighbor-discovery cache](#) on page 3454
[unconfigure netlogin](#) on page 3455
[unconfigure netlogin allowed-refresh-failures](#) on page 3456
[unconfigure netlogin authentication database-order](#) on page 3457
[unconfigure netlogin authentication failure vlan](#) on page 3458
[unconfigure netlogin authentication service-unavailable vlan](#) on page 3458
[unconfigure netlogin banner](#) on page 3459
[unconfigure netlogin dot1x guest-vlan](#) on page 3460
[unconfigure netlogin local-user security-profile](#) on page 3461
[unconfigure netlogin ports](#) on page 3461
[unconfigure netlogin session-refresh](#) on page 3462
[unconfigure netlogin vlan](#) on page 3463
[unconfigure network-clock gptp ports](#) on page 3463
[unconfigure ospf](#) on page 3464
[unconfigure ospfv3](#) on page 3466
[unconfigure pim](#) on page 3467
[unconfigure pim border](#) on page 3468
[unconfigure pim ssm range](#) on page 3469
[unconfigure pim tunnel](#) on page 3470
[unconfigure policy all-rules](#) on page 3471
[unconfigure policy app-signature group name](#) on page 3471
[unconfigure policy autoclear](#) on page 3472
[unconfigure policy captive-portal](#) on page 3473
[unconfigure policy captive-portal listening](#) on page 3474
[unconfigure policy convergence-endpoint all](#) on page 3475
[unconfigure policy convergence-endpoint index](#) on page 3475
[unconfigure policy invalid action](#) on page 3476
[unconfigure policy mactable](#) on page 3477
[unconfigure policy profile](#) on page 3477
[unconfigure policy rule](#) on page 3478
[unconfigure policy syslog](#) on page 3479
[unconfigure policy vlanauthorization](#) on page 3480
[unconfigure port description-string](#) on page 3481
[unconfigure ports display string](#) on page 3482
[unconfigure ports link-flap-detection](#) on page 3482
[unconfigure ports monitor vlan](#) on page 3483
[unconfigure ports redundant](#) on page 3484

[unconfigure process group](#) on page 3485
[unconfigure qosprofile](#) on page 3486
[unconfigure qosprofile wred](#) on page 3487
[unconfigure qosscheduler ports](#) on page 3488
[unconfigure radius](#) on page 3488
[unconfigure radius-accounting](#) on page 3489
[unconfigure radius-accounting server](#) on page 3491
[unconfigure radius server](#) on page 3491
[unconfigure rip](#) on page 3492
[unconfigure ripng](#) on page 3493
[unconfigure sflow](#) on page 3494
[unconfigure sflow agent](#) on page 3495
[unconfigure sflow collector](#) on page 3495
[unconfigure sflow ports](#) on page 3496
[unconfigure slot](#) on page 3497
[unconfigure ssl certificate](#) on page 3498
[unconfigure stacking](#) on page 3499
[unconfigure stacking alternate-ip-address](#) on page 3500
[unconfigure stacking license-level](#) on page 3501
[unconfigure stacking-support](#) on page 3502
[unconfigure stpd ports link-type](#) on page 3503
[unconfigure stpd](#) on page 3504
[unconfigure switch](#) on page 3505
[unconfigure tacacs](#) on page 3506
[unconfigure tacacs-accounting](#) on page 3507
[unconfigure timezone](#) on page 3508
[unconfigure trusted-ports trust-for dhcp-server](#) on page 3508
[unconfigure tunnel](#) on page 3509
[unconfigure upm event](#) on page 3510
[unconfigure upm timer](#) on page 3511
[unconfigure vlan description](#) on page 3512
[unconfigure vlan dhcp](#) on page 3512
[unconfigure vlan dhcp-address-range](#) on page 3513
[unconfigure vlan dhcp-options](#) on page 3514
[unconfigure vlan ipaddress](#) on page 3515
[unconfigure vlan router-discovery](#) on page 3516
[unconfigure vlan router-discovery default-lifetime](#) on page 3517
[unconfigure vlan router-discovery hop-limit](#) on page 3517
[unconfigure vlan router-discovery link-mtu](#) on page 3518
[unconfigure vlan router-discovery managed-config-flag](#) on page 3519
[unconfigure vlan router-discovery max-interval](#) on page 3519
[unconfigure vlan router-discovery min-interval](#) on page 3520

[unconfigure vlan router-discovery other-config-flag](#) on page 3521
[unconfigure vlan router-discovery reachable-time](#) on page 3522
[unconfigure vlan router-discovery retransmit-time](#) on page 3522
[unconfigure vlan subvlan-address-range](#) on page 3523
[unconfigure vlan udp-profile](#) on page 3524
[unconfigure vman ethertype](#) on page 3525
[unconfigure vm-tracking local-vm](#) on page 3526
[unconfigure vm-tracking nms](#) on page 3527
[unconfigure vm-tracking repository](#) on page 3527
[unconfigure vm-tracking vpp vlan-tag](#) on page 3528
[unconfigure vm-tracking vpp](#) on page 3529
[unconfigure vpex](#) on page 3530
[unconfigure vpex mlag-id peer](#) on page 3530
[unconfigure vpls dot1q ethertype](#) on page 3532
[unconfigure vpls snmp-vpn-identifier](#) on page 3533
[unconfigure vr description](#) on page 3533
[unconfigure vr rd](#) on page 3534
[unconfigure vr vpn-id](#) on page 3535
[unconfigure xml-notification](#) on page 3535
[uninstall image](#) on page 3536
[uninstall license file](#) on page 3538
[uninstall license product](#) on page 3539
[upload configuration](#) on page 3541
[upload debug](#) on page 3545
[upload dhcp-bindings](#) on page 3547
[upload log](#) on page 3548
[use configuration](#) on page 3550
[use image](#) on page 3552
[virtual-router](#) on page 3553
[watch](#) on page 3555
[WHILE ... DO](#) on page 3557

The following section explains the operating system commands.

alias

```
alias alias_name command
```

Description

Creates aliases to execute any ExtremeXOS command, including any options, arguments, and redirection.

Syntax Description

<i>alias_name</i>	Specifies an alias name for the command.
<i>command</i>	ExtremeXOS command that you are creating an alias for.

Default

N/A

Usage Guidelines

To be recognized, the alias must be the first word in the string typed at the shell prompt. Substitution does not occur if the alias name string occurs anywhere else. Aliases are only recognized by the EXSH shell session in which they are created.

Executing the command `alias` (with no other arguments) displays a list of current aliases. Executing the command `alias alias_name` displays the command that will be substituted for *alias_name*.

To delete aliases, use the command [unalias](#) on page 3385.

After an alias has been created, you can auto-complete the alias name or display possible aliases along with regular commands by pressing the **TAB** key. You can tab-complete arguments that follow commands corresponding to an alias.

Creating an alias using the name of an existing ExtremeXOS command overrides the original meaning of that command. For example, executing `alias download "download image 102.3.10.5"` allows you to simply type `download image_name` to download your ExtremeXOS image from the 102.3.10.5 location. However, if you then want to download a bootrom file, the command `download bootrom 102.3.10.5 filename` no longer functions correctly. Such an alias can be disabled temporarily and the original command behavior restored by preceding it directly (with no spaces in between) with a backslash, `\download bootrom 102.3.10.5 filename`. This temporarily overrides the alias definition and uses the original command.

To create an alias for a command that contains quoted strings within it, use a backslash. For example, if creating an alias "cr" for the command `configure vlan default description "This is the default VLAN"`, use the command `alias cr "configure vlan default description \"This is the default VLAN\""`.

The following limitations apply to aliases:

- Arguments cannot occur in the middle of alias commands. For example, you cannot create an alias "set_vlan_ip" for the command `configure vlan vlan_name ipaddress ip_address` where you specify the VLAN name as an argument. This is because aliases work through direct textual substitution.
- Aliases cannot be chained together. For example, if you create an alias "sh" for `show version` and another alias "ps" for `process`, then entering `sh ps` at the prompt is not equivalent to entering "show version process".
- You cannot tab-complete commands while trying to create an alias using the `alias` command.
- Aliases cannot be created for the current shell session using UPM scripts or Python scripts.

Aliases are only available in the shell session in which they are created. When you exit the shell your aliases are lost. To create persistent aliases, you need to add the aliases to the script `exshrc.xsf` that you must create using the VI editor and save in the `/usr/local/cfg` folder.

Example

The following example creates an alias named "set" for configure commands:

```
alias set "configure"
```

You can now substitute the command `set` for all configure commands. For example, you can type `set vlan vlan_name tag tag` instead of `configure vlan vlan_name tag tag`.

The following example creates an alias named "mycmd" to substitute for the `configure policy profile` command with the following arguments:

```
alias mycmd "configure policy profile 1 name Extreme pvid 1000 pvid-status enable tci-  
overwrite enable auth-override enable forbidden-vlans 2 cos-status enable cos 2 untagged-  
vlans 2 egress-vlans 200"
```

Typing `mycmd` now executes the command `configure policy profile 1 name Extreme pvid 1000 pvid-status enable tci-overwrite enable auth-override enable forbidden-vlans 2 cos-status enable cos 2 untagged-vlans 2 egress-vlans 200`

The following example lists all current aliases:

```
alias  
alias mycmd='configure policy profile 1 name Extreme pvid 1000 pvid-status enable tci-  
overwrite enable auth-override enable forbidden-vlans 2 cos-status enable cos 2 untagged-  
vlans 2 egress-vlans 200'  
alias set='configure'
```

History

This command was first available in ExtremeXOS 22.3

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

cat

```
cat {--number | -n } {--number-nonblank | -b } {--show-ends | -E } {--  
show-tabs | -T } {--show-nonprinting | -v }
```

Description

Displays the contents of various text files that may be created, edited, or otherwise visible in the user-visible file system.

Syntax Description

--number	Specifies that all output lines are numbered.
-n	Specifies that all output lines are numbered (same as --number).
--number-nonblank	Specifies that nonempty lines be numbers (overrides --number).
-b	Specifies that nonempty lines be numbers (same as --number-nonblank).
--show-ends	Displays \$ at the end of each line.
-E	Displays \$ at the end of each line (same as --show-ends).
--show-nonprinting	Specifies to use ^ and M- notation, except for LF and TAB.
-v	Specifies to use ^ and M- notation, except for LF and TAB (same as --show-nonprinting).
--show-tabs	Displays TAB characters as ^I.
-T	Displays TAB characters as ^I (same as --show-tabs).

Default

N/A

Usage Guidelines

Example files include ExtremeXOS shell script (.xsf), Python (.py), policy (.pol), license (.lic), and logging (.log) files.

There is no restriction on the type of file to be displayed.

CLI paging is effective for this command, but output filtering is not.

Example

The following example displays all options and filetypes for this command:

```
# cat ?
--number          Number all output lines
--number-nonblank Number nonempty output lines (overrides --number)
--show-ends       Display $ at end of each line
--show-nonprinting Use ^ and M- notation, except for LF and TAB
--show-tabs       Display TAB characters as ^I
-E               Display $ at end of each line (same as --show-ends)
-T               Display TAB characters as ^I (same as --show-tabs)
-b               Number nonempty output lines (same as --number-nonblank)
-n               Number all output lines (same as --number)
-v               Use ^ and M- notation, except for LF and TAB (same as --show-
nonprinting)
<filename>       File name
  ../
  VZbase.pol      VZbase1.pol
  VZbase2.pol     cc_logs/
  dhcp-reply.pcap dhcp-reply.pkt
  dhcpreply.pkt   dhcpresp.pkt
  dhcpresp.py     lost+found/
```

```

nsi.dmp
nsi_new.dmp
old_nsi_first_req_after_reboot.dmp
primary.cfg
rest/
test.xsf
nsi_last_req.dmp
nsi_re.dmp
old_nsi_last_req.dmp
rest.tar
ssl/
vmt/

```

History

This command was first available in ExtremeXOS 31.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

cd

```
cd directory_name
```

Description

Changes the current working directory to the directory of the specified file system or relative to the current working directory.

Syntax Description

cd	Change current working directory.
<i>directory_name</i>	Pathname of a directory.

Default

N/A.

Usage Guidelines

Use this command to change the current working directory to the directory of the specified file system.

History

This command was first available in ExtremeXOS 15.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

check policy attribute

```
check policy attribute {attr}
```

Description

Displays the syntax of the specified policy attribute.

Syntax Description

<i>attr</i>	Specifies the attribute check.
-------------	--------------------------------

Default

N/A.

Usage Guidelines

Use this command to display the syntax of policy attributes. The command displays any additional keywords to use with this attribute, and the types of values expected.

Policy attributes are used in the rule entries that make up a policy file.

For each attribute, this command displays which applications use the attribute, and whether the attribute is a match condition or a set (action, action modifier) condition.

The current applications are:

- *ACL (Access Control List)*—access-lists.
- RT—routing profiles, route maps.
- CLF—CLEAR-Flow.

The syntax display does not show the text synonyms for numeric entries. For example, the icmp-type match condition allows you to specify either an integer or a text synonym for the condition. Specifying icmp-type 8 or icmp-type echo-request are equivalent, but the syntax display shows only the numeric option.



Note

The syntax displayed is used by the policy manager to verify the syntax of policy files. The individual applications are responsible for implementing the individual attributes. Inclusion of a particular policy attribute in this command output does not imply that the attribute has been implemented by the application. See the documentation of the particular application for detailed lists of supported attributes.

Example

The following example displays the syntax of the policy attribute icmp-type:

```
check policy attribute icmp-type
```

The following is sample output for this command:

```
( match ) ( ACL )  
icmp-type <uint32 val>
```


History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

check policy

```
check policy policy-name {access-list}
```

Description

Checks the syntax of the specified policy.

Syntax Description

<i>policy-name</i>	Specifies the policy to check.
access-list	Specifies that an access list specific check is performed.

Default

N/A.

Usage Guidelines

Use this command to check the policy syntax before applying it. If any errors are found, the line number and a description of the syntax error are displayed. A policy that contains syntax errors will not be applied.

This command can only determine if the syntax of the policy file is correct and can be loaded into the policy manager database. Since a policy can be used by multiple applications, a particular application may have additional constraints on allowable policies.

Example

The following example checks the syntax of the policy zone5:

```
check policy zone5
```

If no syntax errors are discovered, the following message is displayed:

```
Policy file check successful.
```

History

This command was available in ExtremeXOS 10.1.

The success message and the access-list keyword was added in ExtremeXOS 11.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear access-list counter

```
clear access-list {dynamic} counter {countername} {any | ports port_list  
| vlan vlan_name} {ingress | egress}
```

Description

Clears the specified access list counters.

Syntax Description

dynamic	Specifies that the counter is from a dynamic ACL .
<i>countername</i>	Specifies the ACL counter to clear.
any	Specifies the wildcard ACL.
<i>port_list</i>	Specifies to clear the counters on these ports.
<i>vlan_name</i>	Specifies to clear the counters on the VLAN (Virtual LAN) .
ingress	Clear the ACL counter for packets entering the switch on this interface.
egress	Clear the ACL counter for packets leaving the switch from this interface.

Default

The default direction is ingress; the default ACL type is non-dynamic.

Usage Guidelines

Use this command to clear the ACL counters. If you do not specify an interface, or the any option, you will clear all the counters.

Example

The following example clears all the counters of the ACL on port 2:1:

```
clear access-list counter port 2:1
```

The following example clears the counter counter2 of the ACL on port 2:1

```
clear access-list counter counter2 port 2:1
```

History

This command was first available in ExtremeXOS 10.1.

The **vlan** option was first available in ExtremeXOS 11.0.

The **egress** and **dynamic** options were first available in ExtremeXOS 11.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear access-list meter

```
clear access-list meter {meter_name} [any | ports [all | port_list ] |  
    vlan vlan_name]
```

Description

Clears the specified access list meters.

Syntax Description

<i>meter_name</i>	Specifies the <u>ACL</u> meter to clear.
any	Clear the meter applied to wildcard, including all VLANs and all ports.
ports	Clear the meter applied to a specific port list.
<i>port_list</i>	Specifies to clear the counters on these ports.
vlan	Clear the meter applied to a specific <u>VLAN</u> .
<i>vlan_name</i>	Specifies to clear the counters on the VLAN.

Default

N/A.

Usage Guidelines

Use this command to clear the out-of-profile counters associated with the meter configuration.

Example

The following example clears all the out-of-profile counters for the meters of the ACL on port 2:1:

```
clear access-list meter port 2:1
```

The following example clears the out-of-profile counters for the meter meter2 of the ACL on port 2:1:

```
clear access-list meter meter2 port 2:1
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear account logout

```
clear account [all | name] logout
```

Description

Re-enables an account that has been locked out (disabled) for exceeding the permitted number failed login attempts. This was configured by using the `configure account [all | name] password-policy logout-on-login-failures [on | off]` command.

Syntax Description

all	Specifies all users.
<i>name</i>	Specifies an account name.

Default

N/A.

Usage Guidelines

This command applies to sessions at the console port of the switch as well as all other sessions.

You can re-enable both user and administrative accounts, once they have been disabled for exceeding the 3 failed login attempts.



Note

The failsafe accounts are never locked out.

This command only clears the locked-out (or disabled) condition of the account. The action of locking out accounts following the failed login attempts remains until you turn it off by issuing the `configure account [all | name] password-policy logout-on-login failures off` command.

Example

The following command re-enables the account finance, which had been locked out (disabled) for exceeding 3 consecutive failed login attempts:

```
clear account finance logout
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear bgp flap-statistics

```
clear bgp {neighbor} remoteaddr {address-family [ipv4-unicast |
    ipv4-multicast | ipv6-unicast | ipv6-multicast | vpnv4]} flap-statistics
    [all | rd rd_value | as-path path expression | community [no-advertise
    | no-export | no-export-subconfed | number community_num | AS_Num:Num]
    | network [any / netMaskLen | networkPrefixFilter] {exact}]
```

Description

Clears flap statistics for routes to specified neighbors.

Syntax Description

all	Specifies flap statistics for all routes.
<i>remoteaddr</i>	Specifies the IPv4 or IPv6 address of a <i>BGP</i> (Border Gateway Protocol) neighbor.
ipv4-unicast	Specifies the IPv4 unicast address family.
ipv4-multicast	Specifies an IPv4 multicast address family.
ipv6-unicast	Specifies the IPv6 unicast address family.
ipv6-multicast	Specifies an IPv6 multicast address family.
vpnv4	Specifies the VPNv4 address family for Layer 3 VPN support.
<i>rd_value</i>	Specifies the Route Distinguisher (RD) value for the Layer 3 VPN routes for which you want to clear flap statistics.
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.
<i>community_num</i>	Specifies a community number.
<i>AS_Num</i>	Specifies an autonomous system ID (0-65535).
<i>Num</i>	Specifies a community number.
any	Specifies all routes with a given or larger mask length.
<i>netMaskLen</i>	Specifies a subnet mask length (number of bits).
<i>networkPrefixFilter</i>	Specifies an IP address and netmask.
exact	Specifies an exact match with the IP address and subnet mask.

Default

If no address family is specified, IPv4 unicast is the default.



Note

You must specify an IPv6 address family for an IPv6 peer, because an IPv6 peer does not support the default IPv4 unicast address family. Similarly, if you specify an IPv4 peer and an address family in the command, an IPv4 address family must be specified.

Usage Guidelines

Use this command to clear flap statistics for a specified BGP neighbor.

The option network **any** / *netMaskLen* clears the statistics for all BGP routes whose mask length is equal to or greater than *maskLength*, irrespective of their network address.

The option network **any** / *netMaskLen* exact clears the statistics for all BGP routes whose mask length is exactly equal to *maskLength*, irrespective of their network address.

To clear flap statistics on Layer 3 VPNs, you must configure this feature in the context of the [*MPLS \(Multiprotocol Label Switching\)*](#)-enabled VR; this feature is not supported for BGP routes on the CE (VRF) side of the PE router.

This command applies to the current VR or VRF context.

Example

The following command clears the flap statistics for a specified neighbor:

```
clear bgp neighbor 10.10.10.10 flap-statistics all
```

History

This command was first available in ExtremeXOS 10.1.

The *netMaskLen* options were added in ExtremeXOS 11.0.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 in BGP was added in ExtremeXOS 12.6 BGP.

Support for Layer 3 VPNs was added in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

clear bgp neighbor counters

```
clear bgp neighbor [remoteaddr | all] counters
```

Description

Resets the *BGP* counters for one or all BGP neighbor sessions to zero.

Syntax Description

<i>remoteaddr</i>	Specifies the IPv4 or IPv6 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
- FsmTransitions

The command clear counters also resets all counter for all BGP neighbors. For BGP, the clearcounters command is equivalent to the following BGP command:

```
clear bgp neighbor all counters
```

This command applies to the current VR or VRF context.

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 ExtremeXOS 10.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

clear bootprelay ipv6 prefix-delegation snooping

```
clear bootprelay ipv6 prefix-delegation snooping [ {ipv6-prefix}
    ipv6_prefix |ipv6-prefix all] [ {vlan} vlan_name |vlan all]
```

Description

Clears information about a snooped IPv6 delegate prefix on a VLAN or all VLANs.

Syntax Description

<i>ipv6_prefix</i>	Specifies a snooped IPv6 prefix (/prefix length) delegated via <i>DHCP (Dynamic Host Configuration Protocol)</i> to clear.
ipv6-prefix all	Clears all snooped IPv6 prefixes delegated via DHCP.
<i>vlan_name</i>	Specifies a VLAN.
vlan all	Clears all snooped IPv6 prefixes delegated via DHCP on all VLANs.

Default

N/A

Usage Guidelines

You can clear a specific snooped IPv6 delegated prefix. You can also clear all snooped IPv6 delegated prefixes on a specific VLAN or on all VLANs.

Example

The following example clears information about all snooped IPv6 delegat prefixes on all VLANs.

```
clear bootprelay ipv6 prefix-delegation snooping ipv6-prefix all vlan all
```

History

This command was first available in ExtremeXOS 15.7.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear cdp counters

```
configure cdp counters {ports ports_list}
```

Description

Clears the CDP counter statistics.

Syntax Description

ports	Specifies the ports to clear.
<i>ports_list</i>	Specifies the port list.

Default

N/A.

Usage Guidelines

Use this command to clear the CDP counter statistics.

Example

The following example clears the CDP ports counters:

```
clear cdp counters
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear cdp neighbor

```
clear cdp neighbor [device id device_id | all]
```

Description

Clears the CDP neighbor information.

Syntax Description

device_id	Specifies the Device Identifier to be used in CDP.
<i>device_id</i>	Specifies the Device Identifier of neighbor.
all	Specifies all CDP neighbors.

Default

N/A.

Usage Guidelines

Use this command to clear the CDP neighbor information.

Example

The following command clears all CDP neighbor associations:

```
clear cdp neighbor all
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear counters

```
clear counters
```

Description

Clears all switch statistics and port counters, including port packet statistics, bridging statistics, IP statistics, and log event counters.

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 ports` command to view port statistics. Use the `show log counters` command to show event statistics.

The CLI also provides a number of options that you can specify with the `clear counters` command. If you specify an option, the switch only clears the statistics for that option. For example, if you want to clear, reset only the *STP (Spanning Tree Protocol)* statistics and counters, use the `clear counters stp` command. For more detailed information about those commands, see the specific chapter in the *Switch Engine 32.3 User Guide*.

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.



Note

For the ENTERASYS-POLICY-PROFILE-MIB, the `clear counters` command does not clear counter32.

Example

The following command clears all switch statistics and port counters:

```
clear counters
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear counters bfd

```
clear counters bfd {session | interface}
```

Description

Clears the counters associated with BFD specific settings.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to clear the counters in the BFD session or interface (VLAN). If neither session or interface are specified, the command clears all counters in BFD.

Example

The following command clears all counters in BFD:

```
# clear counters bfd
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear counters bfd missed-hellos

```
clear counters bfd missed-hellos {session-id first {- last} | neighbor
    ipaddress {vr [vrname | all]} | vr [vrname | all]} {current | history
    | both}
```

Description

This command clears the bfd missed hellos counters.

Syntax Description

session-id	Clear counters for sessions having session ID within the given range.
<i>first</i>	Only or first of range of session ID.
<i>last</i>	Last of range of session ID .
neighbor	Neighbor address.
<i>ipaddress</i>	Specify IPv4 or IPv6 destination address.
vr	Virtual router.
<i>vrname</i>	Virtual router name.
all	All virtual routers.
current	Clear only current set of bins.
history	Clear only historical set of bins.
both	Clear both current set and historical set of bins.

Default

Current.

Usage Guidelines

Sessions can be cleared by specifying neighbor IP, by specifying range of session IDs or by specifying VR name. In addition, current bins and historical bins can be cleared separately. These options would help resetting one particular session/bin while tests can run in other sessions/bins.

History

This command was first available in ExtremeXOS 21.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear counters cfm segment all

```
clear counters cfm segment all
```

Description

This command clears both frame-delay and frame-loss information for all existing segments.

Syntax Description

N/A.

Default

N/A.

Usage Guidelines

Use this command to clear both frame-delay and frame-loss information for all existing segments.

Example

```
# clear co cfm seg all
# sho cfm seg
CFM Segment Name      : cs10
Domain Name           : dom1
Association            : a10
MD Level              : 1
Destination MAC       : 00:04:96:52:a7:38
Frame Delay:
DMM Transmission      : In Progress
Transmission Mode     : On Demand
```

```

Total Frames to be sent      : 45
Frames Transmitted          : 0
Pending Frames              : 42
Frames Received             : 0
DMM Tx Interval             : 10 secs
DMR Rx Timeout              : 50 msec
Alarm Threshold              : 10 %
Clear Threshold              : 95 %
Measurement Window Size     : 60
Class of Service            : 6
Tx Start Time               : None
Min Delay                   : None
Max Delay                   : None
Last Alarm Time             : None
Alarm State                 : None
Lost Frames                 : 0
Frame Loss:
LMM Tx Interval             : 10 secs
SES Threshold                : 1.000000e-02
Consecutive Available Count : 4
Measurement Window Size     : 1200
Class of Service            : 6
Total Configured MEPS       : 1
Total Active MEPS           : 1
MEP ID                      : 10
LMM Transmission            : In Progress
Transmission Mode           : On Demand
Total Frames to be sent     : 45
Frames Transmitted          : 0
Pending Frames              : 42
Frames Received             : 0
Availability Status         : Idle
Unavailability Start Time   : None
Unavailability End Time     : None
Tx Start Time               : None
CFM Segment Name           : cs11
Domain Name                 : dom1
Association                  : all
MD Level                    : 1
Destination MAC              : 00:04:96:52:a7:38
Frame Delay:
DMM Transmission            : In Progress
Transmission Mode           : On Demand
Total Frames to be sent     : 45
Frames Transmitted          : 0
Pending Frames              : 42
Frames Received             : 0
DMM Tx Interval             : 10 secs
DMR Rx Timeout              : 50 msec
Alarm Threshold              : 10 %
Clear Threshold              : 95 %
Measurement Window Size     : 60
Class of Service            : 6
Tx Start Time               : Mon Mar 12 10:26:39 2012
Min Delay                   : Mon Mar 12 10:26:49 2012
Max Delay                   : Mon Mar 12 10:26:49 2012
Last Alarm Time             : None
Alarm State                 : None
Lost Frames                 : 0
Frame Loss:
LMM Tx Interval             : 10 secs
SES Threshold                : 1.000000e-02
Consecutive Available Count : 4
Measurement Window Size     : 1200

```

```

Class of Service           : 6
Total Configured MEPs     : 1
Total Active MEPs         : 1
MEP ID                    : 11
LMM Transmission          : In Progress
Transmission Mode         : On Demand
Total Frames to be sent   : 45
Frames Transmitted        : 0
Pending Frames            : 42
Frames Received           : 0
Availability Status       : Idle
Unavailability Start Time : None
Unavailability End Time   : None
Tx Start Time             : None
CFM Segment Name          : cs12
Domain Name               : dom1
Association                : a12
MD Level                  : 1
Destination MAC           : 00:04:96:52:a7:38
Frame Delay:
DMM Transmission          : In Progress
Transmission Mode         : On Demand
Total Frames to be sent   : 45
Frames Transmitted        : 0
Pending Frames            : 42
Frames Received           : 0
DMM Tx Interval           : 10 secs
DMR Rx Timeout            : 50 msec
Alarm Threshold           : 10 %
Clear Threshold           : 95 %
Measurement Window Size   : 60
Class of Service          : 6
Tx Start Time             : Mon Mar 12 10:26:39 2012
Min Delay                 : Mon Mar 12 10:26:49 2012
Max Delay                 : Mon Mar 12 10:26:39 2012
Last Alarm Time           : None
Alarm State               : None
Lost Frames               : 0
Frame Loss:
LMM Tx Interval           : 10 secs
SES Threshold             : 1.000000e-02
Consecutive Available Count : 4
Measurement Window Size   : 1200
Class of Service          : 6
Total Configured MEPs     : 1
Total Active MEPs         : 1
MEP ID                    : 12
LMM Transmission          : In Progress
Transmission Mode         : On Demand
Total Frames to be sent   : 45
Frames Transmitted        : 1
Pending Frames            : 41
Frames Received           : 1
Availability Status       : Available
-----
Total Configured Segments : 11
Total Active Segments     : 11
#
#
#
#

```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear counters cfm segment all frame-delay

```
clear counters cfm segment all frame-delay
```

Description

This command clears only frame-delay information for all existing segments.

Syntax Description

N/A.

Default

N/A.

Usage Guidelines

Use this command to clear only frame-delay information for all existing segments.

Example

```
# clear co cfm seg all frame-delay
#
#
#
# sho cfm segment
CFM Segment Name      : cs10
Domain Name           : dom1
Association            : a10
MD Level              : 1
Destination MAC       : 00:04:96:52:a7:38
Frame Delay:
DMM Transmission      : In Progress
Transmission Mode     : On Demand
Total Frames to be sent : 45
Frames Transmitted     : 1
Pending Frames        : 30
Frames Received       : 1
DMM Tx Interval       : 10 secs
DMR Rx Timeout        : 50 msec
Alarm Threshold        : 10 %
Clear Threshold        : 95 %
Measurement Window Size : 60
Class of Service       : 6
```



```

Tx Start Time           : Mon Mar 12 10:28:59 2012
Min Delay               : Mon Mar 12 10:28:59 2012
Max Delay               : Mon Mar 12 10:28:59 2012
Last Alarm Time         : None
Alarm State             : Not Set
Lost Frames             : 0
Frame Loss:
LMM Tx Interval         : 10 secs
SES Threshold           : 1.000000e-02
Consecutive Available Count : 4
Measurement Window Size : 1200
Class of Service        : 6
Total Configured MEPS   : 1
Total Active MEPS       : 1
MEP ID                 : 10
LMM Transmission        : In Progress
Transmission Mode       : On Demand
Total Frames to be sent : 45
Frames Transmitted      : 4
Pending Frames         : 30
Frames Received         : 4
Availability Status     : Available
Unavailability Start Time : None
Unavailability End Time : None
Tx Start Time           : Mon Mar 12 10:28:29 2012
CFM Segment Name       : cs11
Domain Name             : dom1
Association             : all
MD Level                : 1
Destination MAC         : 00:04:96:52:a7:38
Frame Delay:
DMM Transmission        : In Progress
Transmission Mode       : On Demand
Total Frames to be sent : 45
Frames Transmitted      : 1
Pending Frames         : 30
Frames Received         : 1
DMM Tx Interval         : 10 secs
DMR Rx Timeout          : 50 msec
Alarm Threshold         : 10 %
Clear Threshold         : 95 %
Measurement Window Size : 60
Class of Service        : 6
Tx Start Time           : Mon Mar 12 10:28:59 2012
Min Delay               : Mon Mar 12 10:28:59 2012
Max Delay               : Mon Mar 12 10:28:59 2012
Last Alarm Time         : None
Alarm State             : Not Set
Lost Frames             : 0
Frame Loss:
LMM Tx Interval         : 10 secs
SES Threshold           : 1.000000e-02
Consecutive Available Count : 4
Measurement Window Size : 1200
Class of Service        : 6
Total Configured MEPS   : 1
Total Active MEPS       : 1
MEP ID                 : 11
LMM Transmission        : In Progress
Transmission Mode       : On Demand
Total Frames to be sent : 45
Frames Transmitted      : 12
Pending Frames         : 30
Frames Received         : 12

```

```

Availability Status      : Available
Unavailability Start Time : None
Unavailability End Time  : None
Tx Start Time           : Mon Mar 12 10:27:09 2012
CFM Segment Name        : cs12
Domain Name              : dom1
Association               : a12
MD Level                 : 1
Destination MAC          : 00:04:96:52:a7:38
Frame Delay:
DMM Transmission         : In Progress
Transmission Mode        : On Demand
Total Frames to be sent  : 45
Frames Transmitted        : 1
Pending Frames           : 30
Frames Received           : 1
DMM Tx Interval          : 10 secs
DMR Rx Timeout           : 50 msec
Alarm Threshold           : 10 %
Clear Threshold           : 95 %
Measurement Window Size  : 60
Class of Service          : 6
Tx Start Time            : Mon Mar 12 10:28:59 2012
-----
Total Configured Segments : 11
Total Active Segments     : 11
#
#
#

```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear counters cfm segment all frame-loss

```
clear counters cfm segment all frame-loss
```

Description

This command clears only frame-loss information for all existing segments.

Syntax Description

N/A.

Default

N/A.

Usage Guidelines

Use this command to clear only frame-loss information for all existing segments.

Example

```
# clear co cfm seg all frame-loss
#
#
#
# sho cfm segment
CFM Segment Name      : cs10
Domain Name           : dom1
Association            : a10
MD Level              : 1
Destination MAC        : 00:04:96:52:a7:38
Frame Delay:
DMM Transmission      : In Progress
Transmission Mode     : On Demand
Total Frames to be sent : 45
Frames Transmitted     : 2
Pending Frames        : 29
Frames Received       : 2
DMM Tx Interval       : 10 secs
DMR Rx Timeout        : 50 msec
Alarm Threshold       : 10 %
Clear Threshold       : 95 %
Measurement Window Size : 60
Class of Service      : 6
Tx Start Time         : Mon Mar 12 10:28:59 2012
Min Delay             : Mon Mar 12 10:29:09 2012
Max Delay             : Mon Mar 12 10:29:09 2012
Last Alarm Time       : None
Alarm State           : Not Set
Lost Frames           : 0
Frame Loss:
LMM Tx Interval       : 10 secs
SES Threshold         : 1.000000e-02
Consecutive Available Count : 4
Measurement Window Size : 1200
Class of Service      : 6
Total Configured MEPS : 1
Total Active MEPS     : 1
MEP ID                : 10
LMM Transmission      : In Progress
Transmission Mode     : On Demand
Total Frames to be sent : 45
Frames Transmitted     : 0
Pending Frames        : 29
Frames Received       : 0
Availability Status    : Idle
Unavailability Start Time : None
Unavailability End Time : None
Tx Start Time         : None
CFM Segment Name      : cs11
Domain Name           : dom1
Association            : a11
MD Level              : 1
Destination MAC        : 00:04:96:52:a7:38
Frame Delay:
DMM Transmission      : In Progress
Transmission Mode     : On Demand
```

```

Total Frames to be sent      : 45
Frames Transmitted          : 2
Pending Frames              : 29
Frames Received             : 2
DMM Tx Interval             : 10 secs
DMR Rx Timeout              : 50 msec
Alarm Threshold              : 10 %
Clear Threshold              : 95 %
Measurement Window Size     : 60
Class of Service            : 6
Tx Start Time               : Mon Mar 12 10:28:59 2012
Min Delay                   : Mon Mar 12 10:29:09 2012
Max Delay                   : Mon Mar 12 10:28:59 2012
Last Alarm Time             : None
Alarm State                 : Not Set
Lost Frames                 : 0
Frame Loss:
LMM Tx Interval             : 10 secs
SES Threshold               : 1.000000e-02
Consecutive Available Count : 4
Measurement Window Size     : 1200
Class of Service            : 6
Total Configured MEPs       : 1
Total Active MEPs           : 1
MEP ID                     : 11
LMM Transmission            : In Progress
Transmission Mode           : On Demand
Total Frames to be sent     : 45
Frames Transmitted          : 0
Pending Frames              : 28
Frames Received             : 0
Availability Status         : Idle
Unavailability Start Time   : None
Unavailability End Time     : None
Tx Start Time               : Mon Mar 12 10:29:19 2012
CFM Segment Name           : cs12
Domain Name                 : dom1
Association                  : a12
MD Level                    : 1
Destination MAC              : 00:04:96:52:a7:38
Frame Delay:
DMM Transmission            : In Progress
Transmission Mode           : On Demand
Total Frames to be sent     : 45
Frames Transmitted          : 2
Pending Frames              : 29
Frames Received             : 2
DMM Tx Interval             : 10 secs
DMR Rx Timeout              : 50 msec
Alarm Threshold              : 10 %
Clear Threshold              : 95 %
Measurement Window Size     : 60
Class of Service            : 6
Tx Start Time               : Mon Mar 12 10:28:59 2012
-----
Total Configured Segments    : 11
Total Active Segments       : 11
#
#
#

```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear counters cfm segment frame-delay

```
clear counters cfm segment segment_name frame-delay
```

Description

This command clears only frame-delay information for segment with given segment name.

Syntax Description

<i>segment_name</i>	An alpha numeric string identifying the segment name.
---------------------	---

Default

N/A.

Usage Guidelines

Use this command to clear only frame-delay information for segment with given segment name.

Example

```
# clear co cfm seg cs10 frame-delay
#
#
#
# sho cfm seg cs10
CFM Segment Name      : cs10
Domain Name           : dom1
Association            : a10
MD Level              : 1
Destination MAC        : 00:04:96:52:a7:38
Frame Delay:
DMM Transmission      : In Progress
Transmission Mode     : On Demand
Total Frames to be sent : 45
Frames Transmitted    : 1
Pending Frames        : 34
Frames Received       : 1
DMM Tx Interval       : 10 secs
DMR Rx Timeout        : 50 msec
Alarm Threshold       : 10 %
Clear Threshold        : 95 %
Measurement Window Size : 60
Class of Service       : 6
```

```

Tx Start Time           : Mon Mar 12 10:28:19 2012
Min Delay               : Mon Mar 12 10:28:19 2012
Max Delay               : Mon Mar 12 10:28:19 2012
Last Alarm Time         : None
Alarm State             : Not Set
Lost Frames             : 0
Frame Loss:
LMM Tx Interval         : 10 secs
SES Threshold           : 1.000000e-02
Consecutive Available Count : 4
Measurement Window Size : 1200
Class of Service        : 6
Total Configured MEPs   : 1
Total Active MEPs       : 1
MEP ID                  : 10
LMM Transmission        : In Progress
Transmission Mode       : On Demand
Total Frames to be sent : 45
Frames Transmitted      : 8
Pending Frames          : 34
Frames Received         : 8
Availability Status      : Available
Unavailability Start Time : None
Unavailability End Time  : None
Tx Start Time           : Mon Mar 12 10:27:09 2012
-----
Total Configured Segments : 11
Total Active Segments    : 11
#

```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear counters cfm segment frame-loss mep

```
clear counters cfm segment segment_name frame-loss mep mep_id
```

Description

This command clears only frame-loss information for the given MEP in segment with given segment name.

Syntax Description

<i>segment_name</i>	An alpha numeric string identifying the segment name.
---------------------	---

Default

N/A.

Usage Guidelines

Use this command to clear only frame-loss information for the given MEP in segment with given segment name.

Example

```
# clear counters cfm segment "cs2" frame-loss mep 3
#
#
#
# sho cfm segment
CFM Segment Name           : cs2
Domain Name                 : dom2
Association                  : a2
MD Level                    : 2
Destination MAC              : 00:04:96:52:a7:64
Frame Delay:
DMM Transmission            : Disabled
Frames Transmitted          : 0
Frames Received             : 0
DMM Tx Interval             : 10 secs
DMR Rx Timeout              : 50 msec
Alarm Threshold             : 10 %
Clear Threshold             : 95 %
Measurement Window Size     : 60
Class of Service            : 6
Tx Start Time               : None
Min Delay                   : None
Max Delay                   : None
Last Alarm Time             : None
Alarm State                 : None
Lost Frames                 : 0
Frame Loss:
LMM Tx Interval             : 10 secs
SES Threshold               : 1.000000e-02
Consecutive Available Count : 4
Measurement Window Size     : 1200
Class of Service            : 6
Total Configured MEPs       : 1
Total Active MEPs           : 1
MEP ID                      : 3
LMM Transmission            : In Progress
Transmission Mode           : Continuous
Frames Transmitted          : 0
Frames Received             : 0
Availability Status         : Idle
Unavailability Start Time   : None
Unavailability End Time     : None
Tx Start Time               : None
-----
Total Configured Segments   : 1
Total Active Segments       : 1
#
#
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear counters cfm segment frame-loss

```
clear counters cfm segment segment_name frame-loss
```

Description

This command clears only frame-loss information for segment with given segment name for all associated MEPs.

Syntax Description

<i>segment_name</i>	An alpha numeric string identifying the segment name.
---------------------	---

Default

N/A.

Usage Guidelines

Use this command to clear only frame-loss information for segment with given segment name for all associated MEPs.

Example

```
# clear co cfm seg cs10 frame-loss
#
#
#
# sho cfm seg cs10
CFM Segment Name      : cs10
Domain Name           : dom1
Association            : a10
MD Level              : 1
Destination MAC       : 00:04:96:52:a7:38
Frame Delay:
DMM Transmission      : In Progress
Transmission Mode     : On Demand
Total Frames to be sent : 45
Frames Transmitted    : 1
Pending Frames        : 34
Frames Received       : 1
DMM Tx Interval       : 10 secs
DMR Rx Timeout        : 50 msec
Alarm Threshold       : 10 %
Clear Threshold       : 95 %
Measurement Window Size : 60
Class of Service      : 6
Tx Start Time         : Mon Mar 12 10:28:19 2012
Min Delay             : Mon Mar 12 10:28:19 2012
Max Delay             : Mon Mar 12 10:28:19 2012
```



```

Last Alarm Time      : None
Alarm State          : Not Set
Lost Frames          : 0
Frame Loss:
LMM Tx Interval      : 10 secs
SES Threshold        : 1.000000e-02
Consecutive Available Count : 4
Measurement Window Size : 1200
Class of Service     : 6
Total Configured MEPS : 1
Total Active MEPS    : 1
MEP ID              : 10
LMM Transmission     : In Progress
Transmission Mode    : On Demand
Total Frames to be sent : 45
Frames Transmitted    : 1
Pending Frames       : 33
Frames Received      : 1
Availability Status   : Available
Unavailability Start Time : None
Unavailability End Time : None
Tx Start Time        : Mon Mar 12 10:28:29 2012
-----

```

```

Total Configured Segments : 11
Total Active Segments     : 11
#
#
#
#
#
#

```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear counters cfm segment

```
clear counters cfm segment segment_name
```

Description

This command clears both frame-delay and frame-loss information for segment with given segment name.

Syntax Description

<i>segment_name</i>	An alpha numeric string identifying the segment name.
---------------------	---

Default

N/A.

Usage Guidelines

Use this command to clear both frame-delay and frame-loss information for segment with given segment name.

Example

```
# clear co cfm seg cs2
#
# sho cfm seg cs2
CFM Segment Name      : cs2
Domain Name           : dom1
Association            : a2
MD Level              : 1
Destination MAC        : 00:04:96:52:a7:38
Frame Delay:
DMM Transmission      : In Progress
Transmission Mode     : On Demand
Total Frames to be sent : 45
Frames Transmitted    : 0
Pending Frames       : 40
Frames Received       : 0
DMM Tx Interval       : 10 secs
DMR Rx Timeout        : 50 msec
Alarm Threshold       : 10 %
Clear Threshold       : 95 %
Measurement Window Size : 60
Class of Service      : 6
Tx Start Time         : None
Min Delay             : None
Max Delay             : None
Last Alarm Time       : None
Alarm State           : None
Lost Frames           : 0
Frame Loss:
LMM Tx Interval       : 10 secs
SES Threshold         : 1.000000e-02
Consecutive Available Count : 4
Measurement Window Size : 1200
Class of Service      : 6
Total Configured MEPs : 1
Total Active MEPs     : 1
MEP ID                : 2
LMM Transmission      : In Progress
Transmission Mode     : On Demand
Total Frames to be sent : 45
Frames Transmitted    : 0
Pending Frames       : 40
Frames Received       : 0
Availability Status    : Idle
Unavailability Start Time : None
Unavailability End Time : None
Press <SPACE> to continue or <Q> to quit:
Tx Start Time         : None
-----
Total Configured Segments : 11
Total Active Segments    : 11
#
#
#
#
#
```

```
#
#
#
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear counters cfm session missed-hellos

```
clear counters cfm session missed-hellos { domain_name
    { association_name { ports port_list } { end-point [up|down] } } } }
    { current | history | both }
```

Description

This command clears counters for current or historical cfm session missed-hellos.

Syntax Description

<i>domain_name</i>	IEEE 802.lag Domain name
<i>association_name</i>	IEEE 802.lag or ITU-T Y.1731 Association name
ports	Specify ports to clear counters.
<i>port_list</i>	List of ports to clear counters.
end-point	Specify MEPs (Maintenance association End Point) to clear counters.
up	End point is up.
down	End point is down.
current	Clear only current set of bins.
history	Clear only historical set of bins.
both	Clear both current and historical set of bins.

Default

Current.

Usage Guidelines

None.

History

This command was first available in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear counters edp

```
clear counters edp {ports ports}
```

Description

Clears the counters associated with *EDP (Extreme Discovery Protocol)*.

Syntax Description

<i>ports</i>	Specifies one or more ports or slots and ports.
--------------	---

Default

If you do not specify a port, the EDP counters will be cleared for all ports.

Usage Guidelines

This command clears the following counters for EDP protocol data units (PDUs) sent and received per EDP port:

- Switch PDUs transmitted.
- *VLAN* PDUs transmitted.
- Transmit PDUs with errors.
- Switch PDUs received.
- VLAN PDUs received.
- Received PDUs with errors.

Example

The following command clears the EDP counters on all ports:

```
clear counters edp
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear counters erps

```
clear counters erps ring-name
```

Description

Clear statistics on the specified *ERPS (Ethernet Ring Protection Switching)* ring.

Syntax Description

<i>ring-name</i>	Alphanumeric string that identifies the ERPS ring.
------------------	--

Default

N/A.

Usage Guidelines

Use this command to clear statistics on the specified ERPS ring.

Example

The following command clears statistics on the ERPS ring named “ring1”:

```
clear counters erps ring1
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

clear counters mpls

```
clear counters mpls {[lsp all | [{vlan} vlan_name | vlan all]]}
```

Description

Clears all packet and byte counters for all *MPLS* LSPs and all MPLS protocol counters for all MPLS interfaces.

Syntax Description

lsp all	Clears all MPLS protocol counters for all MPLS LSPs.
<i>vlan_name</i>	Clears all MPLS protocol counters for the MPLS interface on the specified <i>VLAN</i> .
vlan all	Clears all MPLS protocol counters for all MPLS interfaces.

Default

N/A.

Usage Guidelines

This command clears all packet and byte counters for all MPLS LSPs and all MPLS protocol counters for all MPLS interfaces. If the `lsp all` keywords are specified, all packet and byte counters for all MPLS LSPs are cleared. If the `vlan all` keywords are specified, all MPLS protocol counters for all MPLS interfaces are cleared. If a VLAN name is specified, all MPLS protocol counters for the MPLS interface on that VLAN are cleared.

Example

This example clears all MPLS counters associated with VLAN 1:

```
clear counters mpls vlan vlan_1
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support this feature as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

clear counters fdb mac-tracking

```
clear counters fdb mac-tracking [mac_addr | all]
```

Description

Clears the event counters for the *FDB (forwarding database)* MAC-tracking feature.

Syntax Description

<i>mac_addr</i>	Specifies a MAC address, using colon-separated bytes.
all	Clears the counters for all tracked MAC addresses.

Default

N/A.

Usage Guidelines

The `clear counters` command also clears the counters for all tracked MAC addresses.

Example

The following example clears the counters for all entries in the MAC address tracking table:

```
Switch.1 # clear counters fdb mac-tracking all
```

History

This command was first available in ExtremeXOS 12.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear counters flowmon

```
clear counters flowmon
```

Description

This command clears all groups.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

N/A.

Example

The following example clears statistics on all Flow Monitor groups:

```
clear counters flowmon
```

History

This command was first available in ExtremeXOS 32.2.

Platform Availability

This command is available on ExtremeSwitching 5520 and 5420 series switches.

clear counters identity-management

```
clear counters identity-management
```

Description

Clears the identity management feature counters.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command clears the following identity management statistics counters:

- High memory usage level reached count
- Critical memory usage level reached count
- Max memory usage level reached count
- Normal memory usage level trap sent
- High memory usage level trap sent
- Critical memory usage level trap sent
- Max memory usage level trap sent
- Event notification sent

You can view these counters with the `show identity-management statistics` command.



Note

The clear counters command also clears these counters. The following counters relate to active entries and are not cleared: Total number of users logged in, Total number of login instances, and Total memory used.

Example

The following command clears the identity management feature counters:

```
Switch.4 # clear counters identity-management
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear counters iparp

```
clear counters iparp
```

Description

Clears all the IPARP counters.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Example

This example clears all IPARP counters:

```
clear counters iparp
```

History

This command was first available in ExtremeXOS 11.6.

Per virtual router capability was deprecated in ExtremeXOS 30.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear counters l2vpn

```
clear counters l2vpn [vpws [vpws_name | all] | vpls [vpls_name | all]]
```

Description

Clears all the specified VPLS or VPWS counters.

Syntax Description

<i>vpls_name</i>	Identifies the VPLS within the switch (character string).
<i>vpws_name</i>	Identifies the VPWS within the switch (character string).
all	Specifies all VPLS or VPWS VPNs.

Default

N/A.

Usage Guidelines

The l2vpn keyword was introduced in ExtremeXOS Release 12.4 and is required when clearing counters for a VPWS. For backward compatibility, the l2vpn keyword is optional when clearing counters for a VPLS. However, this keyword will be required in a future release, so we recommend that you use this keyword for new configurations and scripts.

Example

This example clears all VPLS counters for the specified VPLS:

```
clear counters vpls myvpls
```

This example clears all VPWS counters for the specified VPWS:

```
clear counters l2vpn vpws myvpws
```

History

This command was first available in ExtremeXOS 11.6.

The l2vpn and vpws keywords were first available in ExtremeXOS 12.4.

Platform Availability

This command is available only on the platforms that support this feature as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

clear counters mpls ldp

```
clear counters mpls ldp {{{vlan} vlan_name} | lsp all}
```

Description

Clears LDP control protocol counters and packet and byte counters associated with LDP LSPs.

Syntax Description

<i>vlan_name</i>	Clears LDP control protocol counters on the specified <i>VLAN</i> .
vlan all	Clears LDP control protocol counters on all <i>MPLS</i> interfaces.
lsp all	Clears all LDP LSP packet and byte counters.

Default

N/A.

Usage Guidelines

By default, all LDP control protocol counters are cleared for all LDP interfaces and all byte counters. Specifying the `vlan` keyword clears only the protocol counters associated with a specified LDP interface. Specifying the `lsp` keyword clears only the packet and byte counters associated with LDP LSPs.

Example

This example clears all LDP control protocol counters and all packet and byte counters for all LDP LSPs:

```
clear counters mpls ldp
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support this feature as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

clear counters mpls rsvp-te

```
clear counters mpls rsvp-te {[lsp all | [{vlan} vlan_name | vlan all]]}
```

Description

Clears all packet and byte counters for all RSVP-TE LSPs and all RSVP-TE protocol counters for all *MPLS* interfaces.

Syntax Description

lsp all	Clears all packet and byte counters for all RSVP-TE LSPs.
<i>vlan_name</i>	Clears all RSVP-TE protocol counters for the MPLS interface on the specified <i>VLAN</i> .
vlan all	Clears all RSVP-TE protocol counters on all MPLS interfaces.

Default

By default, all RSVP-TE control protocol counters are cleared for all RSVP-TE interfaces.

Usage Guidelines

This command clears all packet and byte counters for all RSVP-TE LSPs and all RSVP-TE protocol counters for all MPLS interfaces. If the **lsp all** keywords are specified, all packet and byte counters for all RSVP-TE LSPs are cleared. If the **vlan all** keywords are specified, all RSVP-TE protocol counters for all MPLS interfaces are cleared. If a VLAN name is specified, all RSVP-TE protocol counters for the MPLS interface on that VLAN are cleared.

Example

This example clears the RSVP-TE protocol counters on VLAN 1 only:

```
clear counters mpls rsvp-te vlan vlan_1
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support this feature as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

clear counters mpls static lsp

```
clear counters mpls static lsp {lsp_name | all }
```

Description

Clears the packet and byte counters for one or all static LSPs.

Syntax Description

<i>lsp_name</i>	Identifies the LSP for which counters are to be cleared.
all	Specifies that counters are to be cleared for all static LSPs on this LSR.

Default

N/A.

Usage Guidelines

None.

Example

The following command clears the counters for a static LSP:

```
clear counters mpls static lsp lsp598
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available only on the platforms that support this feature as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

clear counters policy

```
clear counters policy
```

Description

Clears policy rule usage statistics.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command resets the counters on each rule to zero and clears the rule usage.

To see a list of used rules, use the **port-hit** option with the command `show policy rule port-hit {data} {detail} {wide}`.

Example

The following example clears policy rule usage statistics:

```
# clear counters policy
# show policy rule port-hit
No entries found.
```

History

This command was first available in ExtremeXOS 30.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear counters ports

```
clear counters ports {port_list | all}
```

Description

Clears the counters associated with the ports.

Syntax Description

ports	Clears port-related statistics on specified ports or all ports in the system.
<i>port_list</i>	Port list for clear operation.
all	All ports in the system.

Default

All ports.

Usage Guidelines

This command clears the counters for the ports, including the following:

- Statistics.
- Transmit errors.
- Receive errors.
- Collisions.
- Packets.



Note

If you use the clear counters command with no keyword, the system clears the counters for all applications.

Example

The following example clears the counters on all ports:

```
clear counters ports all
```

History

This command was first available in ExtremeXOS 11.3.

This command was updated in ExtremeXOS 15.5 to include the *port_list* variable and the **all** keyword.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear counters ports protocol filter

```
clear counters ports {port_list | all} protocol filter
```

Description

Clears protocol filtering counters.

Syntax Description

<i>port_list</i>	Specifies the port list is separated by a comma (,) or dash (-).
all	Specifies all ports

Default

Disabled.

Usage Guidelines

Use this command to clear protocol filtering counters.

Example

The following example clears all protocol filtering counters:

```
clear counters ports protocol filter
```

The following example clears protocol filtering counters on ports 1-5:

```
clear counters ports 1-5 protocol filter
```

History

This command was first available in ExtremeXOS 15.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear counters stp

```
clear counters stp {[all | diagnostics | domains | ports]}
```

Description

Clears, resets all *STP* statistics and counters.

Syntax Description

all	Specifies all STP domain, port, and diagnostics counters.
diagnostics	Specifies STP diagnostics counters.
domains	Specifies STP domain counters.
ports	Specifies STP port counters.

Default

N/A.

Usage Guidelines

If you do not enter a parameter, the result is the same as specifying the all parameter: the counters for all domains, ports, and diagnostics are reset.

Enter one of the following parameters to reset the STP counters on the switch:

- all—Specifies the counters for all STPDs and ports, and clears all STP counters.
- diagnostics—Clears the internal diagnostic counters.
- domains—Clears the domain level counters.
- ports—Clears the counters for all ports and leaves the domain level counters.

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 that you are monitoring.

Example

The following command clears all of the STP domain, port, and diagnostic counters:

```
clear counters stp
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear counters virtual-network

```
clear counters virtual-network [ all | vn_name ]
```

Description

This command clears statistics (byte/packet counters) on a Virtual Network.

Syntax Description

all	Clear all Virtual Network counters.
<i>vn_name</i>	Clear counters only for the specified Virtual Network string.

Default

N/A.

Usage Guidelines

N/A.

Example

To clear statistics on an existing Virtual Network:

```
clear counters virtual-network vnet44
```

To clear statistics on all Virtual Networks:

```
clear counters virtual-network all
```

History

This command was first available in ExtremeXOS 21.1.

Platform Availability

This command is supported on the ExtremeSwitching 5520 series switch, and stacks with 5420 and 5520 slots only.

clear counters virtual-network remote-endpoint

```
clear counters virtual-network remote-endpoint vxlan [ all | ipaddress
ipaddress]
```

Description

Use this command to clear statistics (byte/packet counters) on a Virtual Network remote endpoint.

Syntax Description

all	Clear all remote endpoint counters.
ipaddress	Clear counters for the specified remote endpoint IP address.
<i>ipaddress</i>	A remote endpoint IP address.

Default

N/A.

Usage Guidelines

N/A.

Example

To clear statistics on an existing Virtual Network remote endpoint:

```
clear counters virtual-network remote-endpoint ipaddress vxlan 10.10.10.146
```

To clear statistics on all Virtual Network remote endpoints:

```
clear counters virtual-network remote-endpoint vxlan all
```

History

This command was first available in ExtremeXOS 21.1.

Platform Availability

This command is supported on the ExtremeSwitching 5520 series switch, and stacks with 5420 and 5520 slots only.

clear counters vpls

```
clear counters vpls [vpls_name | all]
```

Description

Clears all VPLS counters for the specified *vpls_name*.

Syntax Description

<i>vpls_name</i>	Identifies the VPLS within the switch (character string)
all	Specifies all VPLS VPNs.

Default

N/A.

Usage Guidelines

This command clears all VPLS counters for the specified *vpls_name*. If the optional **all** keyword is specified, all packet and byte counters for all VPLS VPNs are cleared.

Example

This example clears all VPLS counters for the specified VPLS:

```
clear counters vpls myvpls
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support this feature as described in in the [Switch Engine 32.3 Feature License Requirements](#) document.

clear counters vr

```
clear counters {vr} vpn-vrf-name
```

Description

Clears statistics information for a VPN Virtual Routing and Forwarding instance (VPN VRF).

Syntax Description

<i>vpn-vrf-name</i>	Specifies the name of a VPN VRF.
---------------------	----------------------------------

Default

N/A.

Usage Guidelines

This command can help to debug control path issues for a VPN VRF. Issuing a global XOS “clear counter” command will also clear VRF counters. This command clears the following counters:

- Route add operation count.
- Route delete operation count.
- Routes dropped count.

This command is supported only on VPN VRFs.

Example

The following command clears the counters for VPN VRF red:

```
Switch.19 # clear counters vr red
```

History

This command was first introduced in XOS Release 15.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear counters vrrp

clear counters vrrp {{**vlan** *vlan_name*} {**vrid** *vridval*}}

Description

Clears, resets all VRRP (Virtual Router Redundancy Protocol) statistics and counters.

Syntax Description

<i>vlan_name</i>	Specifies the name of a VRRP <u>VLAN</u> .
<i>vridval</i>	Specifies the VRRP Router ID (VRID) for a VRRP instance. To display the configured VRRP router instances, enter the <code>show vrrp</code> command.

Default

N/A.

Usage Guidelines

Use this command to reset the VRRP statistics on the switch. Statistics are not reset when you disable and re-enable VRRP.

If you do not enter a parameter, statistics for all VRRP VLANs are cleared.

If you specify only VLAN name, statistics for all VRRP VRIDs on that VLAN are cleared.

If you specify VLAN name and VRRP VRID, only statistics for that particular VRID are cleared.

Example

The following command clears the VRRP statistics on VRRP VLAN v1:

```
clear counters vrrp vlan v1
```

The following command clears the VRRP statistics for VRID 1 on VRRP VLAN v1:

```
clear counters vrrp vlan v1 vrid 1
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the VRRP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

clear counters wred ecn

```
clear counters wred ecn
```

Description

Clears Explicit Congestion Notification (ECN) counters statistics for all ports.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Example

The following example clears ECN counter statistics for all ports:

```
# clear counters wred ecn
```

History

This command was first available in ExtremeXOS 22.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear counters xml-notification

```
clear counters xml-notification {all | target}
```

Description

Clears the statistics counters.

Syntax Description

<i>target</i>	Specifies an alpha numeric string that identifies the configured target.
---------------	--

Default

N/A.

Usage Guidelines

Use this command to unconfigure and reset all statistics counters.

Example

The following command clears all of the xml-notification statistics counters:

```
clear counters xml-notification all
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear cpu-monitoring

```
clear cpu-monitoring {process name}
```

Description

Clears, resets the CPU utilization history and statistics stored in the switch.

Syntax Description

<i>name</i>	Specifies the name of the process.
-------------	------------------------------------

Default

N/A.

Usage Guidelines

When you do not specify any keywords, this command clears the CPU utilization history for the entire switch, including processes, and resets the statistics to zero (0).

When you specify process, the switch clears and resets the CPU utilization history for the specified process.

Example

The following command resets the CPU history and resets statistics to 0 for the TFTP process running on a switch:

```
# clear cpu-monitoring process tftpd
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear dns cache

```
clear dns cache
```

Description

Clears the Domain Name System (DNS) cache entries.

Syntax Description

dns	Domain Name System.
cache	Specifies clearing the DNS cache.

Default

N/A.

Usage Guidelines

None.

Example

The following example clears the DNS cache:

```
# clear dns cache
```

History

This command was first available in ExtremeXOS 30.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear dns cache analytics entries

```
clear dns cache analytics entries {{vr} vr_name}
```

Description

Clears the Domain Name System (DNS) cache analytics entries for a virtual router (VR).

Syntax Description

dns	Domain Name System.
cache	Specifies the DNS cache.
analytics	Specifies the DNS cache analytics.
entries	Specifies clearing the analyzed DNS queries.

vr	Specifies a VR on which to clear entries.
<i>vr_name</i>	Specifies the VR name. If not specified, the VR of the current command context is used.

Default

If not specified, by default the VR of the current command context is used.

Usage Guidelines

This command clears already analyzed DNS queries for a VR. If you do not clear entries with this command, the entries are timed out based on the configured value in the command **configure dns cache analytics** [{**timeout** *minutes*} {**max-entries** *max_entries*}]

Example

The following example clears the DNS cache analytics entries for the current VR:

```
# clear dns cache analytics entries
```

History

This command was first available in ExtremeXOS 30.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear eaps counters

```
clear eaps counters
```

Description

Clears, resets the counters gathered by *EAPS (Extreme Automatic Protection Switching)* for all of the EAPS domains and any EAPS shared ports configured on the switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to clear, reset the EAPS counters.

The counters continue to increment until you clear the information. By clearing the counters, you can see fresh statistics for the time period you are monitoring.

To display information about the EAPS counters, use the following commands:

- `show eaps counters` —This command displays summary EAPS counter information.
- `show eaps counters shared-port` —If configured for EAPS shared ports, this command displays summary EAPS shared port counter information.

Example

The following command clears, resets all of the counters for EAPS:

```
clear eaps counters
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear elrp counters

```
clear elrp counters
```

Description

Clears and resets the ELRP counters.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

You should view the switch statistics before you delete the ELRP counters. Use the `show log counters` command to display event 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 that you are monitoring.

With hard-assisted ELRP, the request to clear ACL counters is sent from ELRP to the ACL manager, and then to hardware one at a time. Since there is one ACL counter per VLAN port, it may take some time for all of the counters to be cleared in hardware when multiple ACL counters are used. If you run the `clear elrp counters` command before all counters are reset, the Pkts-Xmit statistics for some VLANs might temporarily show the sum of partially cleared counters.

Example

The following command clears all switch statistics related to ELRP:

```
# clear elrp counters
```

History

This command was first available in ExtremeXOS 11.1.
Hardware-assisted information was added in ExtremeXOS 30.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear elsm ports auto-restart

```
clear elsm ports port_list auto-restart
```

Description

Clears one or more ELSM-enabled ports that are in the Down-Stuck state.

Syntax Description

<i>port_list</i>	Specifies the ELSM-enabled ports that are permanently in the Down-Stuck state.
------------------	--

Default

N/A.

Usage Guidelines

If you do not have automatic restart enabled, use this command to transition ELSM-enabled ports that are permanently in the Down-Stuck state to the Down state. You can also use the `enable elsm`

`ports port_list auto-restart` command to transition a port from the Down-Stuck state to the Down state.

For information about the ELSM-enabled ports states, see the command `show elsm ports`.

If automatic restart is enabled (this is the default behavior), automatic restart automatically transitions the ports from the Down-Stuck state to the Down state. For more information, see the command `enable elsm ports auto-restart`.

Example

The following command transitions the ports from the Down-Stuck state to the Down state:

```
clear elsm ports 2:1-2:2 auto-restart
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear elsm ports counters

```
clear elsm {ports port_list} counters
```

Description

Clears the statistics gathered by ELSM for the specified ports or for all ports.

Syntax Description

<code>port_list</code>	Specifies the ELSM-enabled ports for which ELSM statistics are being cleared.
------------------------	---

Default

N/A.

Usage Guidelines

You should view the ELSM statistics and counters before you clear them. To view ELSM-specific counter information, use the `show elsm ports all | port_list` command. To view summary ELSM information, including the ports configured for ELSM, use the `show elsm` command.

Use this command to clear only the ELSM-related counters. To clear all of the counters on the switch, including those related to ELSM, use the `clear counters` command.

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 counter, you can see fresh statistics for the time period you are monitoring.

Example

The following command clears the statistics gathered by ELSM for slot 2, ports 1-2:

```
clear elsm ports 2:1-2:2 counters
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear esrp counters

```
clear esrp counters
```

Description

Clears the statistics gathered by *ESRP (Extreme Standby Router Protocol)* for all ESRP domains on the switch.

Syntax Description

This command has no arguments or variables.

Default

None.

Usage Guidelines

Use this command to clear the state transition and the protocol packet counters gathered by ESRP.

The state transition count displays the number of times the ESRP domain entered the following states:

- **Aware**—An Extreme switch that does not participate in ESRP elections but is capable of listening to ESRP Bridge Protocol Data Units (BPDUs).
- **Master**—The master switch is the device with the highest priority based on the election algorithm. The master is responsible for responding to clients for Layer 3 routing and Layer 2 switching for the ESRP domain.

- Neutral—The neutral state is the initial state entered by the switch. In a neutral state, the switch waits for ESRP to initialize and run. A neutral switch does not participate in ESRP elections.
- PreMaster—The pre-master state is an ESRP switch that is ready to be master but is going through possible loop detection prior to transitioning to master.
- Slave—The slave switch participates in ESRP but is not elected or configured the master and does not respond to ARP requests but does exchange ESRP packets with other switches on the same VLAN. The slave switch is available to assume the responsibilities of the master switch if the master becomes unavailable or criteria for ESRP changes.

If the slave is in extended mode, it does not send ESRP hello messages; however, it sends PDUs that can trigger a change in the master switch.

For more information about configuring the ESRP mode of operation on the switch, see the `configure esrp mode [extended | standard]` command. By default, ESRP operates in extended mode.

To display information about the ESRP domain, including the previously described states, use the `show esrp { {name} | {type [vpls-redundancy | standard]} }` command.

The protocol packet count displays the number of times ESRP, ESRP-aware, and ESRP error packets were transmitted and received.

To display information about the ESRP counters, use the `show esrp {name} counters` command.

Example

The following command clears the statistics gathered by ESRP:

```
clear esrp counters
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear esrp neighbor

```
clear esrp esrpDomain neighbor
```

Description

Clears the neighbor information for the specified ESRP domain.

Syntax Description

<i>esrpDomain</i>	Specifies the name of an ESRP domain.
-------------------	---------------------------------------

Default

N/A.

Usage Guidelines

If you add a new switch to your ESRP domain, use this command to clear the existing neighbor information for the ESRP domain. After the switch is up, running, and configured as an ESRP-aware or ESRP-enabled device, new neighbor information is learned.

Before using this command, schedule a downtime for your network. Use this command for maintenance purposes only.

Example

The following example clears the existing neighbor information on the ESRP domain `esrp1` after adding a new switch to the ESRP domain:

```
clear esrp esrp1 neighbor
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear esrp sticky

```
clear esrp esrpDomain sticky
```

Description

Clears the stickiness in the *ESRP* domain and forces the election of the ESRP master switch.

Syntax Description

<i>esrpDomain</i>	Specifies the name of an ESRP domain.
-------------------	---------------------------------------

Default

N/A.

Usage Guidelines

Use the `clear esrp sticky` command to force the election of the ESRP master switch. Before using this command, schedule a downtime for your network.

For example, without stickiness configured, if an event causes the ESRP master to failover to the backup, the previous backup becomes the new master. If another event causes the new master to return to backup, you have experienced two network interruptions. To prevent this, use the `configure esrp election-policy` command and select stickiness as an election algorithm.

If you use sticky as an election metric, and an event causes the ESRP master to failover, ESRP assigns the new master with the highest sticky election metric of 1. Therefore, regardless of changes to the neighbor's election algorithm, the new master retains its position. Sticky is set on the master switch only.

ESRP re-election can occur if sticky is set on the master and a local event occurs. During this time, if the current master has lower election parameters, the backup can become the new master.

If you use `clear esrp esrpDomain sticky` command, it only affects the current master and can trigger ESRP re-election.

Example

The following command clears the stickiness on the ESRP domain esrp1:

```
clear esrp esrp1 sticky
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear ethernet oam counters

```
clear ethernet oam {ports [port_list] counters}
```

Description

Clears Ethernet OAM counters.

Syntax Description

<i>port_list</i>	Specifies the particular port(s).
------------------	-----------------------------------

Default

N/A.

Usage Guidelines

Use this command to clear the Ethernet OAM counters on one or more specified ports. If you do not specify the port(s), counters for all ports are cleared.

Example

The following command clears Ethernet OAM counters on port 2:

```
clear ethernet oam ports 2 counters
```

History

This command was first available in ExtremeXOS 12.3.

Platform Availability

This command is supported on all platforms.

clear fdb

```
clear fdb mac_addr | ports port_list | vlan vlan_name | blackhole | vxlan
       ipaddress remote_ipaddress vr vr_name virtual-network vn_name
```

Description

Clears dynamic *FDB* entries that match the filter.

Syntax Description

<i>mac_addr</i>	Specifies a MAC address, using colon-separated bytes.
<i>port_list</i>	Specifies one or more ports or slots and ports.
<i>vlan_name</i>	Specifies a <i>VLAN</i> name.
blackhole	Specifies the blackhole entries.
vxlan	Specifies VXLAN.
ipaddress	IP address of the remote endpoint.
<i>remote_ipaddress</i>	IPv4 address of the remote tunnel endpoint whose associated FDB entries need to be cleared.
vr	VR/VRF instance the IPv4 address is configured on.
<i>vr_name</i>	An existing VR/VRF name. If not specified, the VR context from where this command is executed is used.

virtual-network	MAC addresses associated with a Virtual Overlay Network learning domain.
<i>vn_name</i>	Name of virtual network whose associated FDB entries need to be cleared.

Default

All dynamic FDB entries are cleared by default.

Usage Guidelines

To clear FDB entries on a given remote endpoint (added to any virtual network):

```
clear fdb vxlan ipaddress remote_ipaddress {vr vr_name}
```

To clear FDB entries on a given remote endpoint added to given virtual network:

```
clear fdb vxlan ipaddress remote_ipaddress {vr vr_name} virtual-network  
vn_name
```

To clear all VXLAN FDB entries (clear all entries learned on the access ports and VXLAN tunnels):

```
clear fdb vxlan
```

Example

The following example clears any FDB entries associated with VLAN corporate:

```
clear fdb vlan corporate
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear fdb vpls

```
clear fdb vpls {vpls_name {peer_ip_address}}
```

Description

Clears the *FDB* information learned for VPLS.

Syntax Description

<i>vpls_name</i>	Clears all FDB entries for the specified VPLS and its associated <u>VLAN</u> .
<i>peer_ip_address</i>	Clears all FDB entries for the pseudowire (PW) associated with the specified VPLS and LDP peer.

Default

N/A.

Usage Guidelines

If the command is used without keywords, every FDB entry learned from any PW is cleared. Using the keywords *vpls_name* clears every FDB entry, (both PW and front panel Ethernet port for the service VLAN) associated with the specified VPLS and the associated VLAN. If the specified VPLS is not bound to a VLAN, the following error message appears:

```
Error: vpls VPLS_NAME not bound to a vlan
```

Using the keywords *vpls_name* and *peer_ip_address* clears all FDB entries from the PW associated with the specified VPLS and LDP peer.

Once the information is cleared from the FDB, any packet destined to a MAC address that has been flushed from the hardware is flooded until the MAC address has been re-learned.

Example

This example clears the FDB information for VPLS 1:

```
clear fdb vpls vpls1
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support this feature as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

clear igmp counters

```
clear igmp counters
```

Description

Clears Internet Group Management Protocol (IGMP) counters.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Example

The following example clears IGMP counters:

```
# clear igmp counters
```

History

This command was first available in ExtremeXOS 30.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IPv4 multicast feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

clear igmp group

```
clear igmp group {grpipaddress} [{vlan} name]
```

Description

Removes one or all *IGMP (Internet Group Management Protocol)* groups.

Syntax Description

<i>grpipaddress</i>	Specifies the group IP address.
<i>name</i>	Specifies a <u>VLAN</u> name.

Default

N/A.

Usage Guidelines

This command can be used by network operations to manually remove learned IGMP group entries instantly. Traffic is impacted until the IGMP groups are relearned. Use this command for diagnostic purposes only.

Example

The following command clears all IGMP groups from VLAN accounting:

```
clear igmp group accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IPv4 multicast feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

clear igmp snooping

```
clear igmp snooping {{vlan} name}
```

Description

Removes one or all [IGMP](#) snooping entries.

Syntax Description

<i>name</i>	Specifies a VLAN name.
-------------	--

Default

N/A.

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 until the snooping entries are learned again.

The dynamic IGMP snooping entries are removed, and then recreated upon the next general query. The static router entry and static group entries are removed and recreated immediately.

This command clears both the IGMPv2 and IGMPv3 snooping entries.

Example

The following command clears IGMP snooping from VLAN accounting:

```
clear igmp snooping accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IGMP snooping feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

clear inline-power stats ports

```
clear inline-power stats ports [all | port_list]
```

Description

Clears the inline statistics for the selected port to zero.

Syntax Description

all	Specifies all ports.
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

N/A.

Usage Guidelines

Use this command to clear all the information displayed by the `show inline-power stats ports port_list` command.

Example

The following command clears the inline statistics for ports 1-8 on a switch:

```
clear inline-power stats ports 1-8
```

The following command displays cleared inline power configuration information for ports 1-8:

```
show inline-power stats ports 1-8
```

Following is sample output from this command:

```
STATISTICS COUNTERS
Port  State      Class    Absent  InvSig  Denied  OverCurrent  Short
1    delivering  class3    0       0       0        0           0
2    delivering  class3    0       0       0        0           0
3    searching   class0    0       0       0        0           0
```

4	searching	class0	0	0	0	0	0
5	searching	class0	0	0	0	0	0
6	searching	class0	0	0	0	0	0
7	searching	class0	0	0	0	0	0
8	searching	class0	0	0	0	0	0

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on the *PoE (Power over Ethernet)* devices listed in *PoE* section of the [Switch Engine 32.3 User Guide](#).

clear ip nat counters vlan

```
clear ip nat counters vlan {vlan_name}
```

Description

Clears the Network Address Translation (NAT) VLAN counters.

Syntax Description

ip	Specifies Internet Protocol (IP).
nat	Specifies NAT.
vlan	Specifies VLAN NAT.
counters	Specifies NAT VLAN counters.
<i>vlan_name</i>	Specifies which VLAN to clear NAT counters for. If no VLAN name is specified, all counters are cleared.

Default

N/A.

Usage Guidelines

To view counter information, run the command `show ip nat vlan counters {vlan_name}`.

Example

The following example clears all NAT VLAN counters:

```
# clear ip nat counters vlan
```

History

This command was first available in ExtremeXOS 31.2.

Platform Availability

This command is available on the switches that support the IP NAT feature. For information about which switches support this and other features, see the [Switch Engine 32.3 Feature License Requirements](#) document.

clear iparp

```
clear iparp {ip_addr {vr vr_name} | vlan vlan_name | vr vr_name}
           {refresh}
```

Description

Removes dynamic entries in the IP ARP table.

Syntax Description

<i>ip_addr</i>	Specifies an IP address.
<i>vlan_name</i>	Specifies a <u>VLAN</u> name.
<i>vr_name</i>	Specifies a Virtual Router (VR) or Virtual Router Forwarding instance (VRF) name.
refresh	Refreshes the ARP cache and deletes the inactive entries.

Default

If you do not specify a VR or VRF, the current VR context is used.

Usage Guidelines

Permanent IP ARP entries are not affected.

This command is specific to a single VR or VRF, and it applies to the current VR context if you do not specify a VR or VRF.

Based on the attributes you specify, the refresh attribute refreshes and deletes the corresponding ARP entries as follows:

- `clear iparp refresh`—Refreshes the entire ARP table and deletes all inactive entries.
- `clear iparp ip_addr refresh`—Refreshes the specified IP address and deletes the IP ARP entry if the ARP request for IP address fails.
- `clear iparp vlan vlan_name refresh`—Refreshes all IP ARP entries associated with the VLAN and deletes all inactive entries for the VLAN.
- `clear iparp vr vr_name refresh`—Refreshes all IP ARP entries associated with the VR and deletes all inactive entries for the VR.

Example

The following example removes a dynamically created entry from the IP ARP table:

```
clear iparp 10.1.1.5
```

The following example refreshes the ARP entry by sending an ARP request for the IP address 10.1.1.5. If the ARP response is received, the dynamic entry is retained; otherwise, the dynamic entry is removed from the IP ARP table if the ARP response is not received.

```
clear iparp 10.1.1.5 refresh
```

History

This command was first available in ExtremeXOS 10.1.

The **refresh** keyword was added in ExtremeXOS 15.7.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear ip-security anomaly-protection notify cache

```
clear ip-security anomaly-protection notify cache {slot [slot | all ]}
```

Description

Clear the local protocol anomaly event cache.

Syntax Description

<i>slot</i>	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

N/A.

Usage Guidelines

This command clears the local protocol anomaly event cache.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear ip-security arp validation violations

```
clear ip-security arp validation violations
```

Description

Clears the violation counters.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command clears the ARP validation violation counters.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear ip-security dhcp-snooping entries

```
clear ip-security dhcp-snooping entries { vlan } vlan_name
```

Description

Clears the DHCP binding entries present on a VLAN.

Syntax Description

<i>vlan_name</i>	Specifies the VLAN of the DHCP server.
------------------	--

Default

N/A.

Usage Guidelines

Use this command to clear the DHCP binding entries present on a VLAN. When an entry is deleted, all its associated entries (such as source IP lockdown, secured ARP, and so on) and their associated ACLs, if any, are also deleted.

Example

The following command clears the DHCP binding entry temporary from the VLAN:

```
clear ip-security dhcp-snooping entries temporary
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear ip-security source-ip-lockdown entries ports

```
clear ip-security source-ip-lockdown entries ports [ ports | all ]
```

Description

Clears locked-down source IP addresses on a per-port basis.

Syntax Description

<i>ports</i>	Specifies the port or ports to be cleared.
all	Specifies that all ports are to be cleared.

Default

N/A.

Usage Guidelines

Use this command to clear locked-down source IP addresses on a per port basis. This command deletes the entries on the indicated ports and clears the associated ACLs.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear ipv6 dad

```
clear ipv6 dad [{vr} vr_name {ipaddress} | vr all | {vlan} vlan_name]
               {counters}
```

Description

Clears the counters for the DAD feature.

Syntax Description

<i>vr_name</i>	Specifies a VR for which to clear the counters.
<i>ipaddress</i>	Specifies an IPv6 address for which to clear the counters.
<i>vlan_name</i>	Specifies a <u>VLAN</u> for which to clear the counters.

Default

If you do not specify a VR or VRF, the command applies to the current VR context.

Usage Guidelines

The vr all option clears the DAD counters for all IPv6 interfaces on the switch.

This command clears the DAD failure counters and removes the MAC for the conflicting IPv6 address after the duplicate address condition has been resolved. The DAD counters and saved MAC addresses are not automatically cleared; they must be cleared with this command.

Example

The following command clears the DAD counters for all IPv6 interfaces in all VRs:

```
clear ipv6 dad vr all
```

History

This command was first available in ExtremeXOS 12.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear isis counters

```
clear isis counters
```

Description

This command clears all IS-IS-related counters in the current virtual router.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command clears all area and VLAN counters.

The following area counters are cleared: corrupted LSPs, LSPDB overloads, manual address from area count, LSP sequence number wraps, LSP sequence number skips, LSP purges, partition changes, and SPF calculations.

The following VLAN counters are cleared: adjacency changes, adjacency initialization failures, rejected adjacencies, ID field length mismatches, maximum area address mismatches, authentication type failures, authentication failures, DIS changes, hello PDU TX and RX count, LSP TX and RX count, CSNP TX and RX count, PSNP TX and RX count, unknown PDU type TX and RX count.

Example

The following command clears all IS-IS counters:

```
clear isis counters
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Premier license as described in the **Switch Engine 32.3 Feature License Requirements** document.

clear isis counters area

```
clear isis counters area [area_name | all]
```

Description

This command clears all IS-IS counters for the specified router process or all router processes.

Syntax Description

<i>area_name</i>	Specifies the router process for which counters are cleared.
all	Clears IS-IS counters for all router processes.

Default

N/A.

Usage Guidelines

The following counters are cleared: corrupted LSPs, LSPDB overloads, manual address from area count, LSP sequence number wraps, LSP sequence number skips, LSP purges, partition changes, SPF calculations, authentication type failures, authentication failures, and ID field length mismatches.

Example

The following command clears the IS-IS counters for areax:

```
clear isis counters area areax
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Premier license as described in the [**Switch Engine 32.3 Feature License Requirements**](#) document.

clear isis counters vlan

```
clear isis counters [vlan all | {vlan} vlan_name]
```

Description

This command clears all IS-IS counters for one or all VLANs.

Syntax Description

vlan all	Clears the counters for all VLANs.
<i>vlan_name</i>	Specifies a single VLAN for which counters are cleared.

Default

N/A.

Usage Guidelines

This command only affects VLANs that have been added to IS-IS router processes. The following counters are cleared: adjacency changes, adjacency initialization failures, rejected adjacencies, ID field length mismatches, maximum area address mismatches, authentication type failures, authentication failures, DIS changes, hello PDU TX and RX count, LSP TX and RX count, CSNP TX and RX count, PSNP TX and RX count, unknown PDU type TX and RX count.

Example

The following command clears the IS-IS counters for all VLANs:

```
clear isis counters vlan all
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Premier license as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

clear l2pt counters rtep

```
clear l2pt counters {[vlan vlan_name [{vxlan{vr vr_name} rtep
    rtep_ipv4}]}
```

Description

Clears L2PT RTEP counters.

Syntax Description

vlan	Optionally clears counters only on a specific VLAN.
<i>vlan_name</i>	Specifies the VLAN name.
vxlan	Specifies Virtual eXtensible LAN.
vr	Specifies Virtual Router.
<i>vr_name</i>	Specifies the Virtual Router Name. If not specified, the VR of the current command context is used.
rtep	Specifies Remote Tunnel End Point.
<i>rtep_ipv4</i>	Specifies the Remote Tunnel End Point IPv4 address.

Default

N/A.

Usage Guidelines

Use this command to clear L2PT RTEP counters.

Example

The following example clears L2PT counters on RTEP 2.2.2.2 of VxLAN interface:

```
clear l2pt counters vlan tenant vxlan rtep 2.2.2.2
```

History

This command was first available in ExtremeXOS 31.3.

Platform Availability

This command is supported on the ExtremeSwitching 5420 and 5520 series switches, and stacks with 5420 and 5520 slots only.

clear l2pt counters vlan

```
clear l2pt counters {vlan vlan_name {ports port_list}}
```

Description

Clears L2PT VLAN counters.

Syntax Description

vlan	Optionally clears counters only on a specific VLAN.
vman	Optionally clears counters only on a specific VMAN.
<i>vlan_name</i>	Specifies the VLAN name.
ports <i>port_list</i>	Optionally clears counters only on specific ports of the VLAN/VMAN. The port list is separated by a comma (,) or dash (-).

Default

Disabled.

Usage Guidelines

Use this command to clear L2PT VLAN counters.

Example

The following example clears all L2PT counters:

```
clear l2pt counters
```

The following example clears L2PT counters on VLAN *vlan1*:

```
clear l2pt counters vlan vlan1
```

History

This command was first available in ExtremeXOS 15.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear l2pt counters vman

```
clear l2pt counters {vman vman_name {ports port_list}}
```

Description

Clears L2PT VMAN counters.

Syntax Description

vlan	Optionally clears counters only on a specific <i>VLAN</i> .
vman	Optionally clears counters only on a specific VMAN.
<i>vlan_name</i>	Specifies the VLAN name.
ports <i>port_list</i>	Optionally clears counters only on specific ports of the VLAN/VMAN. The port list is separated by a comma (,) or dash (-).

Default

Disabled.

Usage Guidelines

Use this command to clear L2PT VMAN counters.

Example

The following example clears all L2PT counters:

```
clear l2pt counters
```

The following example clears L2PT counters on VMAN vlan2:

```
clear l2pt counters vman vlan2
```

History

This command was first available in ExtremeXOS 15.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear l2pt counters vpls

```
clear l2pt counters {[vpls vpls_name {peer ipaddress} | vpws vpws_name]}
```

Description

Clears L2PT counters.

Syntax Description

vpls	Optionally clears counters only on a specific VPLS.
<i>vpls_name</i>	Alpha numeric string identifying VPLS VPN.
peer <i>ipaddress</i>	Optionally clears counters only on a specific peer of the VPLS. The variable specifies an IPv4 address.
vpws <i>vpws_name</i>	Optionally clears counters only on a specific VPWS. The variable is an alphanumeric string identifying the VPWS VPN.

Default

Disabled.

Usage Guidelines

Use this command to clear L2PT counters.

Example

The following example clears L2PT counters on peer 1.1.1.1 of VPLS vpls1:

```
clear l2pt counters vpls vpls1 peer 1.1.1.1
```

History

This command was first available in ExtremeXOS 15.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear lacp counters

```
clear lacp counters
```

Description

Clears the counters associated with Link Aggregations Control Protocol (LACP).

Syntax Description

This command has no parameters or variables.

Default

N/A.

Usage Guidelines

This command clears the following counters for LACP; it sets these counters back to 0 for every LACP port on the device:

- LACP PDUs dropped on non_LACP ports.
- Stats:
 - Rx - Accepted.
 - Rx - Dropped due to error in verifying PDU.
 - Rx - Dropped due to LACP not being up on this port.
 - Rx - Dropped due to matching own MAC.
 - Tx - Sent Successfully.
 - Tx - Transmit error.

Example

The following command clears the LACP counters on all ports:

```
clear lacp counters
```

History

This command was first available in ExtremeXOS 11.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear lldp neighbors

```
clear lldp neighbors [all | port port_list]
```

Description

Clears the LLDP (Link Layer Discovery Protocol) neighbor information collected for one or all ports on the switch.

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
------------------	---

Default

N/A.

Usage Guidelines

LLDP neighbor information for each port is automatically cleared after the period defined by the TTL TLV if no update LLDP protocol data unit (LLDPDU) is received. This command immediately clears the LLDP neighbor information for the specified ports.

Example

The following command clears the LLDP information collected for all ports on the switch:

```
clear lldp neighbors all
```

History

This command was first available in ExtremeXOS 12.4.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear log

```
clear log { static | messages [memory-buffer | nvram] }
```

Description

Clears the log messages in memory and NVRAM.

Syntax Description

static	Specifies that the messages in the NVRAM and memory-buffer targets are cleared.
memory-buffer	Clears entries from the memory buffer.
nvr	Clears entries from NVRAM.

Default

N/A.

Usage Guidelines

The switch log tracks configuration and fault information pertaining to the device.

By default, log entries that are sent to the NVRAM remain in the log after a switch reboot. The `clear log` and `clear log messages memory-buffer` commands remove entries in the memory buffer target; the `clear log static` and `clear log messages nvr` commands remove messages from the NVRAM target. In addition, the `clear log static` command will also clear the memory buffer target.

Execution of these commands on a backup or standby node results in the clearing of that node's information only. Execution of these commands on the master node results in the clearing of information on all nodes in the system.

Example

The following command clears all log messages, from the NVRAM:

```
# clear log static
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear log counters

```
clear log counters [event-condition | [all | event-component]] {severity  
    severity {only}}
```

Description

Clears the incident counters for events.

Syntax Description

<i>event-condition</i>	Specifies the event condition counter to clear.
all	Specifies that all events counters are to be cleared.
<i>event-component</i>	Specifies that all the event counters associated with a particular component should be cleared.
<i>severity</i>	Specifies the minimum severity level of event counters to clear (if the keyword only is omitted).
only	Specifies that only event counters of the specified severity level are to be cleared.

Default

If severity is not specified, then the event counters of any severity are cleared in the specified component.

Usage Guidelines

This command sets the incident counters to zero for each event specified. To display event counters, use the following command: `show log counters`

See the command `show log` for more information about severity levels.

To get a listing of the event conditions in the system, use the following command: `show log events {details}`

To get a listing of the components present in the system, use the following command: `show log components`

In a SummitStack, execution of these commands on a backup or standby node results in the clearing of that node's information only. Execution of these commands on the master node results in the clearing of information on all nodes in the system.

Example

The following example clears the event counters for event conditions of severity error or greater in the component *BGP*:

```
clear log counters "BGP" severity error
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear mac-locking station

```
clear mac-locking station [all | {mac station_mac_address} {first-arrival | static} {ports port_list}]
```

Description

Clears MAC lock station information.

Syntax Description

all	Clears all MAC locking station information for end stations connected to this switch.
<i>station_mac_address</i>	Specifies a MAC address.
first-arrival	Clears first-arrival MAC locking station information.
static	Clears static MAC locking station information.
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

N/A

Usage Guidelines

None.

Example

The following example clears all MAC locking information:

```
clear mac-locking station all
```

History

This command was first available in ExtremeXOS 15.7.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear macsec counters

```
clear macsec counters {ports [port_list]}
```

Description

Clears counters for MAC Security (MACsec) encryption and authentication.

Syntax Description

ports	Specifies port to configure.
<i>port_list</i>	Lists ports to clear MACsec counters on.

Default

Counters for all MACsec ports are cleared unless you choose specific MACsec ports.

Usage Guidelines

This command clears the 4 packet/octet values of the `show macsec ports port-list usage` command, as well as all the statistics shown under the heading “SecY Interface Statistics” of the `show macsec ports port-list detail` command.

Additionally, all MACsec port statistics are cleared by the `clear counters ports {port_list | all}` command.

Example

The following example clears all MACsec counters on port 44:

```
# clear macsec counters ports 44
```

History

This command was first available in ExtremeXOS 30.1.

Platform Availability

This command is available on the following platforms.



Note

The MACsec feature requires the installation of the MAC Security feature pack license.

Platform	Ports
ExtremeSwitching 5320	All ports of all models except stacking ports.
ExtremeSwitching 5420	All ports of all models except stacking ports.
ExtremeSwitching 5520	All ports, except 5520-VIM-4X and 5520-24X 10G ports
ExtremeSwitching 5720	All ports of all models except stacking ports.

clear meter out-of-profile

```
clear meter {metername} out-of-profile {disabled-ports} {status |  
  counters} {ports [ all | portlist | port_group ]}
```


Description

This command allows the clearing of the out-of-profile status and rate-limit counter for meters that have been exceeded. For an input meter and the entered ports, the **status** option will remove the ports from the disabled port out-of-profile list for the entered meter, re-enable ports that may have been disabled due to the out-of-profile meter, and re-enable the syslog and traps for those ports as well. For an input meter and the entered ports, the **counter** option will reset the counters. If neither the status nor counter option is specified, both will be cleared. If the disabled-ports options is specified, only the out-of-profile meters that have disabled ports will be cleared. If no options are specified, all the out-of-profile status and counters will be cleared. If no ports are specified, the command clears the out-of-profile counter for a global meter. Note that the effected counter and status are the aggregates of the rule based counters for both *ACL* and dot1p rules.

Syntax Description

<i>metername</i>	Meter name.
disabled-ports	Clear only the meter out-of-profile status that resulted in disabled-port action.
status	Clear only the meter out-of-profile status.
counters	Clear only the meter counters.
ports	Clear the meter applied to a specified port-list.
all	Clear meter out-of-profile status on all ports.
<i>portlist</i>	Port list separated by a comma or -.
<i>port_group</i>	Port group name.

Default

N/A.

Usage Guidelines

None.

Examples

```
clear meter out-of-profile
clear meter inmeter1 out-of-profile ports 1-5
```

History

This command was first available in ExtremeXOS 16.1.

Platform Availability

This command is available on all platforms.

clear mld counters

```
clear mld counters {{vlan} vlan_name}
```

Description

Clears MLD statistics counters.

Syntax Description

<code>vlan_name</code>	Specifies a VLAN name.
------------------------	--

Default

N/A.

Usage Guidelines

Use this command to manually clear MLD statistics counters.

Example

The following example clears all MLD counters for all VLANs:

```
clear mld counters
```

If a VLAN is specified, only the counters on the specific VLAN is cleared.

History

This command was first available in ExtremeXOS 15.2.

Platform Availability

This command is available on the platforms listed for the IPv6 multicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

clear mld group

```
clear mld group {v6grpipaddress} {{vlan} name}
```

Description

Removes one or all MLD groups.

Syntax Description

<i>v6grpipaddress</i>	Specifies the group IP address.
<i>name</i>	Specifies a <u>VLAN</u> name.

Default

N/A.

Usage Guidelines

This command is used to manually remove learned MLD group entries instantly.

Example

The following command clears all MLD groups from VLAN accounting:

```
clear mld group accounting
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 multicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

clear mld snooping

```
clear mld snooping {{vlan} name}
```

Description

Removes one or all MLD snooping entries.

Syntax Description

<i>name</i>	Specifies a <u>VLAN</u> name.
-------------	-------------------------------

Default

N/A.

Usage Guidelines

This command can be used by network operations to manually remove MLD snooping entries instantly. However, removing an MLD snooping entry can disrupt the normal forwarding of multicast traffic until the snooping entries are learned again.

The static and dynamic MLD snooping entries are removed, and then recreated upon the next general query. The static router entry is removed and recreated immediately.

Example

The following command clears MLD snooping from VLAN accounting:

```
clear mld snooping accounting
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 multicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

clear msdp counters

```
clear msdp counters {peer remoteaddr | peer all | system} {vr vrname}
```

Description

This command resets the *MSDP (Multicast Source Discovery Protocol)* counters to zero.

Syntax Description

peer all	Specifies all MSDP peers.
<i>remoteaddr</i>	Specifies the IP address of the MSDP peer.
system	Clears the global MSDP counters.
<i>vrname</i>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

N/A.

Usage Guidelines

The `clear msdp counters` command clears the following MSDP counters:

- Per peer counters:
 - Number of SA messages received.
 - Number of SA messages transmitted.
 - Number of SA request messages received.
 - Number of SA request messages transmitted.
 - Number of SA response messages received.
 - Number of SA response messages transmitted.
 - Number of SA messages received without encapsulated data.
 - Number of SA messages transmitted without encapsulated data.
 - Number of SA messages received with encapsulated data.
 - Number of SA messages transmitted with encapsulated data.
 - Number of times the MSDP peer attained an “ESTABLISHED” state.
 - Number of times the peer-RPF check failed.
 - Number of times the TCP connection attempt failed.
 - Total number of received messages.
 - Total number of transmitted messages.
- Global counters:
 - None defined.

The `clear counters` command will also clear all MSDP counters, but it clears the counters for all other applications too.

Example

The following command clears the counters for an MSDP peer with the IP address 192.168.45.43:

```
clear msdp counters peer 192.168.45.43
```

The following command clears the all peer and global counters:

```
clear msdp counters
```

The following command clears all counters for a particular peer:

```
clear msdp counters peer 192.168.32.45
```

The following command clears the counters of all MSDP peers:

```
clear msdp counters peer all
```

The following command clears the global counters:

```
clear msdp counters system
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

clear msdp sa-cache

```
clear msdp sa-cache {{peer} remoteaddr | peer all} {group-address grp-addr} {vr vrname}
```

Description

This command purges all SA cache entries and notifies the PIM that the SA cache is empty.

Syntax Description

peer all	Specifies all <i>MSDP</i> peers. All matching SA cache entries from all peers are removed from the database.
<i>grp-addr</i>	Specifies the IP address and subnet mask of the multicast group you want to clear. All SA cache entries that match the specified group address are removed from the database.
<i>remoteaddr</i>	Specifies the IP address of the MSDP peer. All matching SA cache entries learned from the specified peer are removed from the database.
<i>vrname</i>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

N/A.

Usage Guidelines

MSDP receives SA messages periodically. After clearing SA cache entries from the local database, MSDP relearns those entries during the next advertisement from its peer.

Example

The following example clears SA cache records for an MSDP peer with the IP address 192.168.45.43:

```
clear msdp sa-cache peer 192.168.45.43
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

clear neighbor-discovery cache

```
clear neighbor-discovery cache ipv6 {ipv6address {vr vr_name} | vlan
vlan_name | vr vr_name} refresh
```

Description

Deletes a dynamic entry from the neighbor cache.

Syntax Description

<i>vr_name</i>	Specifies a VR or VRF.
<i>ipv6address</i>	Specifies an IPv6 address.
<i>vlan_name</i>	Specifies an IPv6 configured <u>VLAN</u> .
refresh	Refreshes the IPv6 neighbor discovery cache and deletes the inactive entries.

Default

N/A.

Usage Guidelines

This command clears dynamic entries from the neighbor cache. The **vr** option is used to specify the VR or VRF on which the operation is performed. When this option is omitted, it applies to current VR context.

When the **ipv6address** or **vlan** options are specified, only the entries with matching IPv6 addresses or that correspond to that VLAN are cleared.

Based on the attributes you specify, the refresh attribute refreshes and deletes the corresponding IPv6 neighbor discovery entries as follows:

- `clear neighbor-discovery cache refresh`—Refreshes the entire IPv6 neighbor discovery cache and deletes all inactive entries.
- `clear neighbor-discovery cache ipv6address refresh`—Refreshes the specified neighbor-discovery entry and deletes the neighbor-discovery entry if the neighbor solicitation for the IP address fails.
- `clear neighbor-discovery cache vlan vlan_name refresh`—Refreshes all neighbor-discovery entries associated with the VLAN and deletes all inactive entries for the VLAN.
- `clear neighbor-discovery cache vr vr_name refresh`—Refreshes all neighbor-discovery entries associated with the VR and deletes all inactive entries for the VR.

Example

The following example clears all entries from the neighbor cache:

```
clear neighbor-discovery cache
```

The following example refreshes all entries in the neighbor discovery cache and delete inactive entries if the neighbor solicitation fails:

```
clear neighbor-discovery cache refresh
```

History

This command was first available in ExtremeXOS 11.2.

The **refresh** option was added in ExtremeXOS 15.7.1.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

clear netlogin state

```
clear netlogin state {port port_list}
```

Description

Clears and initializes the network login sessions on a VLAN port.

Syntax Description

<i>port_list</i>	Specifies the ports to clear.
------------------	-------------------------------

Default

None.

Usage Guidelines

Clear the states of every MAC learned on this VLAN port and put the port back to unauthenticated state. The port will be moved to its original VLAN if configured in campus mode.

Example

The following command clears the Network Login state of port 2:9:

```
clear netlogin state port 2:9
```


History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear netlogin state agent

```
clear netlogin state agent portportlist [dot1x |mac |web-based]
```

Description

Clears the *NetLogin* authentication state.

Syntax Description

port <i>portlist</i>	Clears only for the specified ports.
dot1x	Clears only the 802.1x authentication state.
mac	Clears only the MAC authentication state.
web-based	Clears only web-based authentication state.

Default

N/A

Example

The following example clears the dot1x authentication state on port 1:

```
clear netlogin state agent port 1 dot1x
```

History

This command was first available in ExtremeXOS 21.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear netlogin state mac-address

```
clear netlogin state mac-address mac
```

Description

Initialize/reset the network login sessions for a specified supplicant.

Syntax Description

<i>mac</i>	Specifies the MAC address of the supplicant.
------------	--

Default

N/A.

Usage Guidelines

This command is essentially equivalent to a particular supplicant logging out. The MAC address will be cleared from the *FDB*, the port is put back to its original *VLAN* (for campus mode), and the port state is set to unauthenticated, if this was the last authenticated MAC on this port.

Example

The following command resets the Network Login session for the supplicant with the MAC address of 00:e0:18:01:32:1f:

```
clear netlogin state mac-address 00:e0:18:01:32:1f
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear network-clock gtp counters

```
clear network-clock gtp ports counters {ports [port_list | all]}
```

Description

Clears gTP port counters.

Syntax Description

gtp	IEEE 802.1AS Generalized Precision Time Protocol.
counters	gTP port counters.

<i>port_list</i>	Specifies one or more of the switch's physical ports.
all	Specifies all of the switch's physical ports.

Default

N/A.

Usage Guidelines

Use this command to clear gTP port counters. The command `clear counters` also clears the gTP port counters (along with all other counters).

Example

```
clear network-clock gtp counters
clear network-clock gtp counters ports 2-4
clear network-clock gtp counters ports all
```

History

This command was first available in ExtremeXOS 15.3

Platform Availability

This command is available on all platforms if the AVB feature pack license is installed on the switch.

clear nodealias

```
clear nodealias { ports [port_list | all] | alias-id alias_id }
```

Description

This command clears alias entries out of the Node Alias feature database. You can clear information by specified port(s) or alias ID. Node Alias discovers information about the end systems on a per-port basis. Information from packets from end systems, such as VLANID, source MAC address, source IP address, protocol, etc. are captured in a database that can be queried.

Syntax Description

nodealias	Node Alias feature that maps source IP address, MAC address, host name, and protocol on a per port basis.
ports	Designates that you want to clear node alias information for the selected ports.
<i>port_list</i>	Specifies from which ports to clear node alias information. Designated as a port list separated by comma (,) or dash (-).
all	Clears node alias information from all ports.

alias-id	Designates that you want to clear node alias information for the specified alias ID from all ports.
<i>alias_id</i>	Specifies the alias ID that you want information cleared for from the database.

Default

None.

Usage Guidelines

If the port is part of a LAG (Link Aggregation Group), this command is only allowed on the master port.

Example

The following example clears all node alias entries on port 7:

```
clear nodealias ports 7
```

The following example clear node alias entries for alias ID 716168949 from all ports:

```
clear nodealias alias-id 716168949
```

History

This command was first available in ExtremeXOS 22.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear ospf counters

```
clear ospf counters { interfaces [all | vlan vlan_name | area area-identifier] | area [all | area-identifier] | virtual-link [all | router-identifier area-identifier] | neighbor [all | routerid [ip-address {ip-mask} | ipNetmask] | vlan vlan_name] | system}
```

Description

Clears the OSPF (Open Shortest Path First) counters (statistics).

Syntax Description

<i>vlan_name</i>	Specifies a <u>VLAN</u> name.
<i>router-identifier</i>	Specifies a router interface number.
<i>area-identifier</i>	Specifies an OSPF area.

<i>ip-address</i>	Specifies an IP address
<i>ip-mask</i>	Specifies a subnet mask.
<i>ipNetmask</i>	Specifies IP address / Netmask.
system	Specifies the OSPF system counters.

Default

N/A.

Usage Guidelines

The global command clear counters also clears all OSPF counters. This global command is the equivalent of clear ospf counters for OSPF.

Example

The following command clears the OSPF counters for area 1.1.1.1:

```
clear ospf counters area 1.1.1.1
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

clear ospfv3 counters

```
clear ospfv3 counters {interfaces [[vlan | tunnel] all | {vlan} vlan-name | {tunnel} tunnel-name | area area-identifier] | virtual-link
[all | {routerid} router-identifier {area} area-identifier]}
```

Description

Clears the [OSPFv3 \(Open Shortest Path First version 3\)](#) counters (statistics).

Syntax Description

all	Specifies all VLANs , tunnels, areas, neighbors, or virtual-links.
<i>vlan_name</i>	Specifies an IPv6 configured VLAN.
<i>tunnel_name</i>	Specifies an IPv6 tunnel.

<i>router-identifier</i>	Specifies a router identifier, a four-byte, dotted decimal number.
<i>area_identifier</i>	Specifies an OSPFv3 area, a four-byte, dotted decimal number.

Default

N/A.

Usage Guidelines

The global command clear counters also clears all OSPFv3 counters. This global command is the equivalent of clear ospfv3 counters for OSPFv3.

This command can be used to clear various OSPFv3 counters (Interface, Area, Virtual-Link, System etc.). The following is the list of various counters that would be reset to zero by this command:

- Neighbor specific counters:
 - Number of state changes.
 - Number of events.
- Interface/VLAN/Virtual-link/Tunnel specific counters:
 - Hellos Rxed
 - Hellos Txed
 - DB Description Rxed
 - DB Description Txed
 - LSA Request Rxed
 - LSA Request Txed
 - LSA Update Rxed
 - LSA Update Txed
 - LSA Ack Rxed
 - LSA Ack Txed
 - In Discards

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [**Switch Engine 32.3 Feature License Requirements**](#) document.

clear pim cache

```
clear pim {ipv4 | ipv6} cache {group_addr {source_addr}}
```

Description

Clears PIM multicast cache table.

Syntax Description

ipv4	Specifies an IPv4 address.
ipv6	Specifies an IPv6 address.
<i>group_addr</i>	Specifies a group address.
<i>source_addr</i>	Specifies a source IP address.

Default

If no options are specified, all PIM cache entries are flushed.

Usage Guidelines

This command can be used by network operators to manually remove IPMC software and hardware forwarding cache entries instantly. If the stream is available, caches are re-created; otherwise, caches are removed permanently. This command can disrupt the normal forwarding of multicast traffic.

Example

The following example resets the IP multicast table for group 224.1.2.3:

```
clear pim cache 224.1.2.3
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

clear pim snooping

```
clear pim snooping {vlan} name
```

Description

Clears all PIM snooping neighbors, joins received on the VLAN, and the VLAN forwarding entries.

Syntax Description

<i>name</i>	Specifies the VLAN to which this command applies.
-------------	---

Default

N/A.

Usage Guidelines

None.

Example

The following command clears the PIM snooping database for the Default VLAN:

```
clear pim snooping "Default"
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM snooping feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

clear port rate-limit flood

```
clear port [all | port_list | port_group] rate-limit flood out-of-  
profile {disabled-ports} {status | counter}
```

Description

This command clears the counter and/or status of ports of a flood rate-limiter that may have had their limit exceeded.

Syntax Description

<i>port_list</i>	Clears a port list.
<i>port_group</i>	Clears a port group.
all	Clears all ports.
out-of-profile	Clears only out-of-profile rate-limiters.
disabled-ports	Clears only ports that have been disabled due to out-of-profile status.
both	Clears out-of-profile status and counter for rate-limiter.

status	Clears only out-of-profile status for rate-limiter.
counter	Clear only out-of-profile counter for rate-limiter.

Default

All.

Usage Guidelines

The `clear ports rate-limit flood out-of-profile` command allows the clearing of the counter and/or status of ports of a flood rate-limiter that may have had their limit exceeded. For the entered ports, the status option removes the ports from the disabled port out-of-profile list, re-enables ports that may have been disabled due to out-of-profile rate-limit, and re-enables the syslog and traps for those ports as well. For the entered ports, the counter option resets the counters. If neither option is specified, both the status and counter will be cleared. If the **disabled-ports** option is specified, only the out-of-profile statuses that have disabled ports will be cleared. If no options are specified, all out-of-profile statuses will be cleared.

Example

```
clear ports all rate-limit flood out-of-profile
clear ports all rate-limit flood out-of-profile disabled-ports
clear ports fldGroupA rate-limit flood out-of-profile status
clear ports 1-24 rate-limit flood out-of-profile counter
```

History

This command was first available in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear ports link-flap-detection counters

```
clear ports [port_list | all] link-flap-detection counters
```

Description

Clears the counters related to port link-flapping.

Syntax Description

ports	Physical ports.
<i>port_list</i>	List of ports that you want to clear the link flap counters on.

all	Selects all ports in the system to have their link-flap counters cleared.
counters	Counters related to link flapping.

Default

N/A

Example

The following example clears the link flap counters for ports 4 through 12:

```
clear ports 4-12 link-flap-detection counters
```

History

This command was first available in ExtremeXOS 22.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear ports link-flap-detection status

```
clear ports [port_list | all] link-flap-detection status
```

Description

Manually enables ports that have been disabled due to excessive link-flapping.

Syntax Description

ports	Physical ports.
<i>port_list</i>	List of ports that you want to enable that were disabled due to excessive link flapping.
all	Enables all ports in the system that were disabled due to excessive link flapping.
status	Enable ports currently in disabled state due to excessive link flapping.

Default

N/A

Usage

Ports that have been disabled due to excessive link flapping cannot be enabled using the `enable port` command. They must be enabled using the `clear ports link-flap-detection status` command.

Example

The following example re-enables all ports on the switch that were disabled due to excessive link flapping:

```
clear ports all link-flap-detection status
```

History

This command was first available in ExtremeXOS 22.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear port rate-limit flood

```
clear port [all | port_list | port_group] rate-limit flood out-of-  
profile {disabled-ports} {status | counter}
```

Description

This command clears the counter and/or status of ports of a flood rate-limiter that may have had their limit exceeded.

Syntax Description

<i>port_list</i>	Clears a port list.
<i>port_group</i>	Clears a port group.
all	Clears all ports.
out-of-profile	Clears only out-of-profile rate-limiters.
disabled-ports	Clears only ports that have been disabled due to out-of-profile status.
both	Clears out-of-profile status and counter for rate-limiter.
status	Clears only out-of-profile status for rate-limiter.
counter	Clear only out-of-profile counter for rate-limiter.

Default

All.

Usage Guidelines

The `clear ports rate-limit flood out-of-profile` command allows the clearing of the counter and/or status of ports of a flood rate-limiter that may have had their limit exceeded. For the entered ports, the status option removes the ports from the disabled port out-of-profile list, re-enables ports that may have been disabled due to out-of-profile rate-limit, and re-enables the syslog and traps for those ports as well. For the entered ports, the counter option resets the counters. If neither option is specified, both the status and counter will be cleared. If the **disabled-ports** option is specified, only the out-of-profile statuses that have disabled ports will be cleared. If no options are specified, all out-of-profile statuses will be cleared.

Example

```
clear ports all rate-limit flood out-of-profile
clear ports all rate-limit flood out-of-profile disabled-ports
clear ports fldGroupA rate-limit flood out-of-profile status
clear ports 1-24 rate-limit flood out-of-profile counter
```

History

This command was first available in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear process group statistics

```
clear process group statistics {vital | other}
```

Description

This command clears the memory- and CPU-related statistics of "Vital" and/or "Other" (non-vital) groups.

Syntax Description

statistics	Designates clearing statistics for the process groups.
vital	Selects clearing statistics for the "vital" process group. If you make no selection, statistics for both groups are cleared.
other	Selects clearing statistics for the "non-vital" process group. If you make no selection, statistics for both groups are cleared.

Default

If you make no selection, statistics for both groups are cleared.

Example

The following example clears statistics for the "vital" group:

```
clear process group statistics vital
```

History

This command was first available in ExtremeXOS 22.2.

The **exos** option was removed in ExtremeXOS 31.5.

The **vital** option was first available in ExtremeXOS 31.5

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear rip counters

```
clear rip counters
```

Description

Clears the *RIP (Routing Information Protocol)* counters (statistics).

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command clears the RIP statistics counters:

```
# clear rip counters
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms with a Base license or higher as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

clear ripng counters

```
clear ripng counters {vlan vlan-name | tunnel tunnel-name}
```

Description

Clears the [RIPng \(Routing Information Protocol Next Generation\)](#) global or interface-specific counters (statistics).

Syntax Description

<i>vlan-name</i>	Specifies an IPv6 configured VLAN .
<i>tunnel-name</i>	Specifies an IPv6 tunnel.

Default

N/A.

Usage Guidelines

None.

Example

The following command clears the RIPng statistics counters:

```
clear ripng counters
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Edge, Advanced Edge, or Core license. For licensing information, see the [Switch Engine 32.3 Feature License Requirements](#) document.

clear screen

```
clear screen
```

Description

This command clears the screen of a login session with the termcaps-defined capability and returns the prompt to the top.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

History

This command was first available in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear session

```
clear session [history | sessId | all]
```

Description

Terminates a Telnet and/or SSH2 sessions from the switch.

Syntax Description

history	Clears the chronology of sessions that were opened.
<i>sessId</i>	Specifies a session number from show session output to terminate.
all	Terminates all sessions.

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 output of this command displays information about current Telnet and/or SSH2 sessions including:

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

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.

When invoked to clear the session history, the command clears the information about all the previous sessions that were logged. The information about the active sessions remains intact.

Example

The following example terminates session 4 from the system:

```
clear session 4
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear slot

```
clear slot slot
```

Description

Clears a slot of a previously assigned module type.

Syntax Description

<code>slot</code>	Specifies the slot number.
-------------------	----------------------------

Default

N/A.

Usage Guidelines

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

If a node is configured for one type of switch, and a different type of switch is inserted in the stack, the inserted node is put into a mismatch state (where the inserted node does not match the configured node), and is not brought online. To use the new switch type in a node, the node configuration must be cleared or configured for the new switch type. Use the enable mirroring to port tagged command to configure the node.

Example

The following command clears node 2 of a previously assigned switch type:

```
clear slot 2
```

The following command clears slot 4 of a previously assigned switch type in a stack:

```
clear slot 4
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on SummitStacks.

clear snmp notification-log

```
clear snmp notification-log [counters | entries] { default | name hex  
                                hex_name}
```

Description

Clears entries and counters from a notification log.

Syntax Description

counters	Specifies to clear counters.
entries	Specifies to clear notification entries.
default	Optionally clear just the default log.
<i>name</i>	Optionally clear just the specified log.
hex	Provide value in hexadecimal.
<i>hex_name</i>	Optionally clear just the specified log (log name in hexadecimal).

Default

Disabled.

Usage Guidelines

Use this command to clear entries and counters from a notification log.

Example

The following example clears global counters:

```
clear snmp notification-log counters
```

The following example clears all entries from all logs:

```
clear snmp notification-log entries
```

The following example clears counters for the default log:

```
clear snmp notification-log counters default
```

The following example clears all entries from nmslog1:

```
clear snmp notification-log entries nmslog1
```

History

This command was first available in ExtremeXOS 15.5.

The **default** and **hex** keywords and *hex_name* variable were added in ExtremeXOS 15.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear stpd ports

```
clear stpd stpd_name ports port_list protocol-migration
```

Description

Resets the partner Spanning Tree Protocol version to the configured version.

Syntax Description

<i>stpd_name</i>	Specifies an <i>STPD (Spanning Tree Domain)</i> name on the switch.
<i>port_list</i>	Specifies the port list, which can be separated with a comma or a dash.
protocol-migration	Resets the partner protocol mode to configured mode.

Default

N/A

Usage Guidelines

STP detects the spanning tree version on a network and sends out the equivalent BPDU. If this switch receives a legacy IEEE 802.1D configuration BPDU (a BPDU with the protocol version set to 0), Protocol Migration feature supports the forcefully allowing the user to choose the version, where a switch supporting *MSTP (Multiple Spanning Tree Protocol)* is forced to behave as STP or RSTP.

For example, three bridges on shared media, two of are configured dot1w (RSTP) and one is dot1d (legacy STP) mode

These bridges will transmit STP BPDUs on their connected ports since one of the peers is in dot1d mode. If the dot1d mode configured bridge leaves this shared media the remaining two bridges will keep sending STP BPDUs even though they should use RTP BPDUs normally. By using this feature we can clear the STP BPDU transmission and starts sending the RSTP BPDUs.

Example

The following example resets the protocol migration for the port 1:10 in STP domain r1:

```
clear stpd r1 ports 1:10 protocol-migration
```

History

This command was first available in ExtremeXOS 15.7.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear switch bluetooth

```
clear switch bluetooth device [all | address]
```

Description

Clears either all paired Bluetooth devices or a particular paired device.

Syntax Description

switch	Designates clearing switch information.
bluetooth	Designates clearing Bluetooth information.
device	Designates clearing Bluetooth devices.
all	Clears all Bluetooth devices.
<i>address</i>	Clears only the Bluetooth device at the specified MAC address.

Default

N/A.

Usage Guidelines

To clear all paired Bluetooth devices, use the **all** option.

To clear only a specific device, use the *address* option. To find the address of a specific Bluetooth device, use the show **switch bluetooth** [**statistics** | **inventory**] command without the **statistics** option.

To enable Bluetooth capabilities, use the enable **switch bluetooth** {**discovery** | **pairing**} command.

Example

The following example clears all Bluetooth devices:

```
# clear switch bluetooth device all
```

The following example clears the Bluetooth device at address 00:04:96:9a:46:48:

```
# clear switch bluetooth device 00:04:96:9a:46:48
```

History

This command was first available in ExtremeXOS 30.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

clear vm storage

```
clear vm storage
```

Description

Formats the virtual machine (VM) storage module (SSD) for use.

Syntax Description

vm	Designates a virtual machine.
storage	Specifies formatting disk storage (VM storage module) for use by VMs.

Default

N/A.

Usage Guidelines

None.

Example

The following example formats VM storage:

```
# clear vm storage
```

History

This command was first available in ExtremeXOS 30.3.

Platform Availability

This command is available on all platforms that support the Integrated Application Hosting (IAH) feature and have a Core or Premier license installed. For a list of platforms that support the IAH feature and for information about licenses, see the [Switch Engine 32.3 Feature License Requirements](#).

clear vlan dhcp-address-allocation

```
clear vlan vlan_name dhcp-address-allocation [[all {offered | assigned |  
declined | expired}] | ipaddress]
```

Description

Removes addresses from the [DHCP](#) allocation table.

Syntax Description

<i>vlan_name</i>	Specifies the VLAN of the DHCP server.
all	Specifies all IP addresses, or all IP addresses in a particular state.
offered	Specifies IP addresses offered to clients.
assigned	Specifies IP addresses offered to and accepted by clients.
declined	Specifies IP addresses declined by clients.
expired	Specifies IP addresses whose lease has expired and not renewed by the DHCP server.
<i>ipaddress</i>	Specifies a particular IP address.

Default

N/A.

Usage Guidelines

You can delete either a single entry, using the IP address, or all entries. If you use the all option, you can additionally delete entries in a specific state.

Example

The following command removes all the declined IP addresses by hosts on the VLAN temporary:

```
clear vlan temporary dhcp-address-allocation all declined
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure access-list

```
configure access-list aclname [any | ports port_list | vlan vlan_name]
                        {ingress | egress}
```

Description

Configures an access list to the specified interface.

Syntax Description

<i>aclname</i>	Specifies the <u>ACL</u> policy file name.
any	Specifies that this ACL is applied to all interfaces as the lowest precedence ACL.
<i>port_list</i>	Specifies the ingress or egress port list on which the ACL is applied.
<i>vlan_name</i>	Specifies the <u>VLAN</u> on which the ACL is applied.
ingress	Apply the ACL to packets entering the switch on this interface.
egress	Apply the ACL to packets leaving the switch from this interface. (ExtremeSwitching X460-G2, X670-G2, X440-G2, X465, X620 series switches only).

Default

The default direction is ingress.

Usage Guidelines

The access list applied in this command is contained in a text file created either externally to the switch or using the [edit policy](#) command. The file is transferred to the switch using TFTP before it is applied to the ports. The ACL name is the file name without its ".pol" extension. For example, the ACL blocknetfour would be in the file blocknetfour.pol.

Specifying the keyword any applies the ACL to all the ports, and is referred to as the wildcard ACL. This ACL is evaluated for ports without a specific ACL applied to it, and is also applied to packets that do not match the ACL applied to the interface.

Example

The following command configures the ACL policy test to port 1:2 at ingress:

```
configure access-list test ports 1:2
```

The following command configures the ACL mydefault as the wildcard ACL:

```
configure access-list mydefault any
```

The following command configures the ACL policy border as the wildcard egress ACL:

```
configure access-list border any egress
```

History

This command was first available in ExtremeXOS 10.1.

The VLAN option was first available in ExtremeXOS 11.0.

The egress option was first available in ExtremeXOS 11.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure access-list action-resolution highest-priority

```
configure access-list action-resolution highest-priority
```

Description

This command puts user [ACLs](#) into "highest priority only" action resolution mode.

Syntax Description

This command has no arguments or variables.

Default

Multiple.

Usage Guidelines

Use this command to put user ACLs into "highest priority only" action resolution mode. All of the static policies and dynamic ACL rules that are installed after this command has been executed execute only the actions of the highest priority rule that has being matched, even if there are matches in the lower priority virtual slices with non-conflicting actions. This behavior is achieved by putting all virtual slices used by user ACLs into the same virtual group. However, all the policies and dynamic ACL rules that were installed prior to the execution of this command would stay in their separate virtual groups. As a result of this, the rules installed prior to the execution of this command will execute non- conflicting actions from the matches in lower priority virtual slices in addition to executing all the actions of the highest priority match. If a save and reboot was done after this command has being executed, all static policies and dynamic ACL rules will operate in "highest priority only" action resolution mode.

History

This command was first available in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure access-list action-resolution multiple

```
configure access-list action-resolution multiple
```

Description

This command puts user ACLs into "multiple matches" action resolution mode. All the static policies and dynamic ACL rules that are installed after this command is entered would execute all the actions of the highest priority rule that has being matched as well as all non conflicting actions from the matches in the lower priority virtual slices.

Syntax Description

This command has no arguments or variables.

Default

Multiple.

Usage Guidelines

Use this command to put user ACLs into "multiple matches" action resolution mode. All the static policies and dynamic ACL rules that are installed after this command is entered would execute all the

actions of the highest priority rule that has been matched as well as all non-conflicting actions from the matches in the lower priority virtual slices.

This behavior is achieved by putting all virtual slices used by user ACLs into separate virtual groups. However, all the policies and dynamic ACL rules that were installed prior to the execution of this command would stay in their old single virtual group. As a result, the rules installed prior to the execution of this command will execute only the actions of the highest priority match. If the save and reboot was done after this command has being executed, all static policies and dynamic ACL rules will operate in "multiple matches" action resolution mode. "Multiple matches" is the default mode on the switch, and if none of action-resolution commands has being executed the switch will operate in "multiple matches" resolution mode.

History

This command was first available in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure access-list add

```
configure access-list add dynamic_rule [ [[first | last] {priority
    p_number} {zone zone} ] | [[before | after] rule] | [ priority
    p_number {zone zone} ]] [ any | vlan vlan_name | ports port_list ]
    {ingress | egress}
```

Description

Configures a dynamic ACL rule to the specified interface and sets the priority and zone for the ACL.

Syntax Description

<i>dynamic_rule</i>	Specifies a dynamic ACL rule.
first	Specifies that the new dynamic rule is to be added as the first rule.
last	Specifies that the new dynamic rule is to be added as the last rule.
priority	Priority of rule within a zone.
<i>p_number</i>	Specifies the priority number of the rule within a zone. The range is from 0 (highest priority) to 7 (lowest priority).
<i>zone</i>	Specifies the ACL zone for the rule.
before rule	Specifies that the new dynamic rule is to be added before an existing dynamic rule.
after rule	Specifies that the new dynamic rule is to be added after an existing dynamic rule.
any	Specifies that this ACL is applied to all interfaces.
<i>vlan_name</i>	Specifies the <u>VLAN</u> on which this ACL is applied.

<i>port_list</i>	Specifies the ports on which this ACL is applied.
ingress	Apply the ACL to packets entering the switch on this interface.
egress	Apply the ACL to packets leaving the switch from this interface.

Default

The default direction is ingress.

Usage Guidelines

The dynamic rule must first be created before it can be applied to an interface. Use the following command to create a dynamic rule:

```
create access-list dynamic-rule conditions actions {non-permanent}
```

When a dynamic ACL rule is applied to an interface, you will specify its precedence among any previously applied dynamic ACLs. All dynamic ACLs have a higher precedence than any ACLs applied through ACL policy files.

Specifying the keyword *any* applies the ACL to all the ports, and is referred to as the wildcard ACL. This ACL is evaluated for ports without a specific ACL applied to them, and is also applied to packets that do not match the ACL applied to the interface.

The *priority* keyword can be used to specify a sub-zone within an application's space. For example, to place ACLs into three sub-zones within the CLI application, you can use three priority numbers, such as 2, 4, and 7.

Configuring priority number 1 is the same as configuring first priority. Configuring priority number 8 is the same as configuring last priority.

Example

The following command applies the dynamic ACL *icmp-echo* as the first (highest precedence) dynamic ACL to port 1:2 at ingress:

```
configure access-list add icmp-echo first ports 1:2
```

The following command applies the dynamic ACL *udpacl* to port 1:2, with a higher precedence than rule *icmp-echo*:

```
configure access-list add udpacl before icmp-echo ports 1:2
```

History

This command was first available in ExtremeXOS 11.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure access-list delete

```
configure access-list delete ruleName [ any | vlan vlan_name | ports
port_list | all] {ingress | egress}
```

Description

Removes a dynamic [ACL](#) rule from the specified interface.

Syntax Description

<i>ruleName</i>	Specifies a dynamic ACL rule name.
any	Deletes this ACL as the wildcard ACL.
<i>vlan_name</i>	Specifies the VLAN on which this ACL is deleted.
<i>port_list</i>	Specifies the ports on which this ACL is deleted.
all	Deletes this ACL from all interfaces.
ingress	Deletes the ACL for packets entering the switch on this interface.
egress	Deletes the ACL for packets leaving the switch from this interface.

Default

The default direction is ingress.

Usage Guidelines

Specifying the keyword **all** removes the ACL from all interfaces it is used on.

Example

The following command removes the dynamic ACL icmp-echo from the port 1:2:

```
configure access-list delete icmp-echo ports 1:2
```

History

This command was first available in ExtremeXOS 11.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure access-list network-zone

```
configure access-list network-zone zone_name [add | delete] [mac-address
  macaddress {macmask} | ipaddress [ipaddress {netmask} | ipNetmask |
  ipv6_address_mask]]
```

Description

Adds or removes IP and MAC addresses to and from the network-zone.

Syntax Description

network-zone	Logical group of remote devices.
<i>zone_name</i>	Specifies the network-zone name.
add	Adds a logical group of entities to the network-zone.
delete	Deletes a logical group of entities to the network-zone.
mac-address	MAC address.
<i>macaddress</i>	Specifies the MAC address to be added/removed to/from the network-zone.
<i>macmask</i>	Specifies the MAC Mask. Example FF:FF:FF:00:00:00.
ipaddress	Specifies IPv4 address.
<i>ipaddress</i>	Specifies the IP address.
<i>netmask</i>	Specifies IP netmask.
<i>ipNetmask</i>	Specifies the IP address/Netmask.
<i>ipv6_address_mask</i>	Specifies IPv6 address/IPv6 prefix length.

Default

N/A.

Usage Guidelines

Use this command to to add or remove IP/MAC addresses to/from the network-zone.

Example

The following command adds an IPv6 IP address to network-zone “zone1”:

```
Switch# configure access-list network-zone zone1 add ipaddress
11.1.1.1/32
```

If you try to add the same IP/MAC with the same or narrow mask, the configuration is rejected, with the following error message.

```
Switch #configure access-list network-zone "zone1" add ipaddress 11.1.1.1/24
```

```
Error: Network Zone "zone1" - Zone already has the same entity value with same or wider mask.
```

If you try to add more than eight attributes to a network-zone, the following error message is printed.

```
Switch #configure access-list network-zone "zone1" add ipaddress 11.1.1.1/24
Error: Network Zone "zone1" - Reached maximum number of attributes. Unable to add more.
```

History

This command was first available in ExtremeXOS 15.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure access-list rule-compression port-counters

```
configure access-list rule-compression port-counters [shared |
dedicated]
```

Description

Switches between *ACL* configuration modes.

Syntax Description

shared	Sharing is “on” for counter rules.
dedicated	Sharing is “off” for counter rules.

Default

Dedicated.

Usage Guidelines

Use this command to switch between two ACL configuration modes. In the first mode, “port-counters shared”, similar port-based ACL rules with counters are allowed to share the same hardware entry. This uses less space but provides an inaccurate counter value. In the second mode, “port-counters dedicated”, similar port-based ACL rules with counters are not allowed to share the same hardware entry, thereby consuming more entries but providing a precise count.

Only ACLs that are entered after this command is entered are affected. The command does not affect any ACLs that are already configured.

To configure all ACLs in shared mode, configure access-list rule-compression port-counters shared must be entered before any ACLs are configured or have been saved in the configuration when a switch is booted.

This is a global setting for the switch; that is, the option does not support setting some ACL rules with shared counters and some with dedicated counters.

To view the results of the configuration use the [show access-list configuration](#) command.

Example

The following command configures ACL rules with counters to share the same hardware entry:

```
configure access-list rule-compression port-counters shared
```

History

This command was first available in ExtremeXOS 12.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure access-list vlan-acl-precedence

```
configure access-list vlan-acl-precedence [dedicated | shared]
```

Description

Configures precedence mode for policy-file based ACLs that are applied on a VLAN.

Syntax Description

dedicated	Allocates exclusive precedence for VLAN-based ACLs.
shared	VLAN-based ACLs share the precedence with other ACLs.

Default

Dedicated.

Usage Guidelines

The following feature applies to only policy-file based ACLs that are applied on a VLAN. Use this command to switch between two VLAN-based ACL configuration modes. In the shared vlan-aclprecedence mode, VLAN-based ACL rules share the same precedence with other types of ACL rules and provides the same behavior as in the previous software releases. In the dedicated vlan-acl-precedence mode, VLAN-based ACL rules have different precedence compared to other types of ACL rules and this is the default mode. The dedicated mode yields improved installation performance for VLAN based access-lists but may affect hardware rule utilization in some configurations.

After configuring, you are prompted to reboot the system for the changes to take effect.

Example

The following command allocates exclusive precedence for VLAN-based static ACL rules:

```
configure access-list vlan-acl-precedence dedicated
```

History

This command was first available in ExtremeXOS 12.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure access-list width

```
configure access-list width [double | single] [slot slotNo | all]
```

Description

Configures the TCAM width of a switch.

Syntax Description

double	Specifies a double wide <i>ACL</i> TCAM. Provides double wide ACL key with additional qualifiers.
single	Specifies a single wide ACL TCAM.
<i>slotNo</i>	Specifies the slot to configure.
all	Specifies all slots.

Default

Single.

Usage Guidelines



Note

This command is not applicable to the ExtremeSwitching X870 series switches. Key width is applied automatically on X870 switches.

Use this feature to configure the width of the ACL TCAM key of a slot or switch to be either double wide or single wide.

The switch must be rebooted for the configuration change to take effect.

If you attempt to configure a double wide mode on a slot or switch that does not support it, an error message is displayed.

To display the configured mode, use the [show access-list width](#) command.

Example

The following command configures slot 1 to use double wide mode:

```
# configure access-list width double slot 1
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure access-list zone

```
configure access-list zone name zone-priority number
```

```
configure access-list zone name move-application appl_name to-zone name  
application-priority number
```

```
configure access-list zone name {add} application appl_name  
application_priority number
```

```
configure access-list zone name delete application appl_name
```

Description

Configures the priority of a zone; moves an application from one zone to another at a specified priority; adds an application to a zone with a specified priority, or changes the priority of an application within a zone; deletes an application from a zone.

Syntax Description

<i>name</i>	Specifies a zone name.
zone-priority <i>number</i>	Sets the priority of the zone.
move-application <i>appl_name</i>	Specifies the name of an application to be moved.
to-zone <i>name</i>	Specifies the zone to which the application is moved.
application-priority <i>number</i>	Sets the priority of the application within the zone. The range is from 0 (highest priority) to 7 (lowest priority).
add	Adds an application to a zone at a specified priority.
application <i>appl_name</i>	Specifies the application to be added to the zone.
application_priority <i>number</i> <i>number</i>	Sets the priority of a new or existing application within a zone. The range is from 0 (highest priority) to 7 (lowest priority).

Default

N/A.

Usage Guidelines

To configure the priority of a specific zone, use the syntax:

```
configure access-list zone name zone-priority number
```

To move an application from one zone to another, and set its priority in the new zone, use the syntax:

```
configure access-list zone name move-application appl-name to-zone name  
application-priority number
```

To add an application to a zone and specify its priority or to change the priority of an application within a zone, use the syntax:

```
configure access-list zone name {add} application appl-name  
application_priority number
```

To delete an application from a zone, use the syntax:

```
configure access-list zone name delete application appl-name
```

Example

The following command adds the CLI application to the zone myzone at a priority of 6:

```
configure access-list zone myzone add cli application-priority 6
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure account

```
configure account [all | name]
```

Description

Configures a password for the specified account, either user account or administrative account.

Syntax Description

all	Specifies all accounts (and future users).
<i>name</i>	Specifies an account name.

Default

N/A.

Usage Guidelines

You must create a user or administrative account before you can configure that account with a password.

Use the `create account` command to create a user account.

The system prompts you to specify a password after you enter this command. You must enter a password for this command; passwords cannot be null and cannot include the following characters: "<", ">", and "?".



Note

Once you issue this command, you cannot have a null password. However, if you want to have a null password (that is, no password on the specified account), use the `create account` command.

Passwords can have a minimum of 0 character and can have a maximum of 32 characters. Passwords are case-sensitive. User names are not case-sensitive.



Note

If the account is configured to require a specific password format, the minimum is 8 characters. See `configure account password-policy char-validation` for more information.

You must have administrator privileges to change passwords for accounts other than your own.

Example

The following example defines a new password green for the account marketing:

```
configure account marketing
```

The switch responds with a password prompt:

```
password: green
```

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: green
```

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

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure account encrypted

```
configure account [all | name] encrypted e-password
```

Description

Encrypts the password that is entered in plain text for the specified account, either user account or administrative account.

Syntax Description

all	Specifies all accounts (and future users).
<i>name</i>	Specifies an account name.
<i>e-password</i>	Enter in plain text the string you for an encrypted password. See Usage Guidelines for more information.

Default

N/A.

Usage Guidelines

You must create a user or administrative account before you can configure that account with a password.

Use the [create account](#) account command to create a user account.

When you use this command, the following password that you specify in plain text is entered and displayed by the switch in an encrypted format. Administrators should enter the password in plain text. The encrypted password is then used by the switch once it encrypts the plain text password. The encrypted command should be used by the switch only to show, store, and load a system-generated encrypted password in configuration; this applies with the following commands: [save configuration](#), [show configuration](#), and [use configuration](#).



Note

Once you issue this command, you cannot have a null password. However, if you want to have a null password (that is, no password on the specified account), use the [create account](#) command.

Passwords can have a minimum of 0 character and can have a maximum of 32 characters. Passwords are case-sensitive. User names are not case-sensitive.



Note

If the account is configured to require a specific password format, the minimum is 8 characters. See [configure account password-policy char-validation](#) for more information.

You must have administrator privileges to change passwords for accounts other than your own.

Example

The following command encrypts the password red for the account marketing:

```
configure account marketing encrypted red
```

History

This command was first available in ExtremeXOS 11.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure account password-policy char-validation

```
configure account [all | name] password-policy char-validation [none | all-char-groups]
```

Description

Requires that the user include an upper-case letter, a lower-case letter, a digit, and a symbol in the password.

Syntax Description

all	Specifies all users (and future users).
<i>name</i>	Specifies an account name.
none	Resets password to accept all formats.
all-char-groups	Specifies that the password must contain at least two characters from each of the four groups. Note: The password minimum length will be eight characters if you specify this option.

Default

N/A.

Usage Guidelines

This feature is disabled by default.

Once you issue this command, each password must include at least two characters of each of the following four types:

- Upper-case A-Z.
- Lower-case a-z.
- 0-9.
- !, @, #, \$, %, ^, *, (,).

The minimum number of characters for these specifically formatted passwords is 8 characters and the maximum is 32 characters.

Use the none option to reset the password to accept all formats.

Example

The following example requires all users to use this specified format for all passwords:

```
configure account all password-policy char-validation all-char-groups
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure account password-policy history

```
configure account [all | name] password-policy history [num_passwords |  
none]
```

Description

Configures the switch to verify the specified number of previous passwords for the account. The user is prevented from changing the password on a user or administrative account to any of these previously saved passwords.

Syntax Description

all	Specifies all accounts (and future users).
<i>name</i>	Specifies an account name.
<i>num_passwords</i>	Specifies the number of previous passwords the system verifies for each account. The range is 1 to 10 passwords.
none	Resets the system to not remember any previous passwords.

Default

N/A.

Usage Guidelines

Use this command to instruct the system to verify new passwords against a list of all previously used passwords, once an account successfully changes a password.

The limit is the number of previous passwords that the system checks against in the record to verify the new password.

If this parameter is configured, the system returns an error message if a user attempts to change the password to one that is saved by the system (up to the configured limit) for that account; this applies to both user and administrative accounts. This also applies to a configured password on the default admin account on the switch.

The limit of previous passwords that the system checks for previous use is configurable from 1 to 10. Using the none option disables previous password tracking and returns the system to the default state of no record of previous passwords.

Example

The following command instructs the system to verify that the new password has not been used as a password in the previous 5 passwords for the account engineering:

```
configure account engineering password-policy history 5
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure account password-policy logout-on-login-failures

```
configure account [all | name] password-policy logout-on-login-failures
[on | off]
```

Description

Disables an account after the user has three consecutive failed login attempts.

Syntax Description

all	Specifies all users (and future users).
<i>name</i>	Specifies an account name.
on	Specifies an account name.
off	Resets the password to never lockout the user.

Default

N/A.

Usage Guidelines

If you are not working on SSH, you can configure the number of failed logins that trigger lockout, using the `configure cli max-failed-logins num-of-logins` command.

This command applies to sessions at the console port of the switch as well as all other sessions and to user-level and administrator-level accounts. This command locks out the user after 3 consecutive failed login attempts; the user's account must be specifically re-enabled by an administrator.

Using the off option resets the account to allow innumerable consecutive failed login attempts, which is the system default. The system default is that three failed consecutive login attempts terminate the particular session, but the user may launch another session; there is no lockout feature by default.



Note

The switch does not allow to lock out of at least one administrator account.

Example

The following command enables the account finance for lockout.

After three consecutive failed login attempts, the account is subsequently locked out:

```
configure account finance password-policy lockout-on-login-failures on
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure account password-policy lockout-time-period

```
configure account [all | name] password-policy lockout-time-period  
[num_mins | until-cleared]
```

Description

This command allows you to configure the lockout time period (ranging from one minute to one hour).

Syntax Description

all	Configure all accounts.
<i>name</i>	Configure a specific account name.
<i>num_min</i>	Number of minutes (1-60) account is locked after max-failed-logins, unless unlocked via <code>clear account <i>name</i> lockout</code> .
until-cleared	Account is locked after max-failed-logins until unlocked via <code>clear account <i>name</i> lockout</code> .

Default

Until-cleared.

Usage Guidelines

Use this command to configure the lockout time period (ranging from one minute to one hour. Note that fail safe and admin accounts will also be locked out if lockout time period is specified. If there is more than one admin account, admin will be locked out even if the lockout time period is set to indefinite.

History

This command was first available in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure account password-policy max-age

```
configure account [all | name] password-policy max-age [num_days | none]
```

Description

Configures a time limit for the passwords for specified accounts. The passwords for the default admin account and the failsafe account do not age out.

Syntax Description

all	Specifies all accounts (and future users).
<i>name</i>	Specifies an account name.
<i>num_days</i>	Specifies the length of time that a password can be used. The range is 1 to 365 days.
none	Resets the password to never expire.

Default

N/A.

Usage Guidelines

The passwords for the default admin account and the failsafe account never expire.

The time limit is specified in days, from 1 to 365 days. Existing sessions are not closed when the time limit expires; it will not open the next time the user attempts to log in.

When a user logs into an account with an expired password, the system first verifies that the entered password had been valid prior to expiring, and then prompts the user to change the password.



Note

This is the sole time that a user with a user-level (opposed to an administrator-level) account can make any changes to the user-level account.

Using the **none** option prevents the password for the specified account from ever expiring (it resets the password to the system default of no time limit).

To set a minimum lifespan for passwords, use the `configure account [all | name] password-policy min-age [num_days | none]` command.

In the case of conflicting settings between these two commands, a setting requiring a password change overrides a setting that prohibits a password change. For example, if **max-age** is set to 10 days, thus requiring a password change in 10 days, and a **min-age** is set to 20 days, attempting to forbid a password change until 20 days, the configuration to change the password after 10 days takes precedence over the configuration to not change the password for 20 days.

To view the current selection for the maximum lifespan for passwords, use the `show accounts password-policy` command.

Example

The following command sets a 3-month time limit for the password for the account marketing:

```
# configure account marketing password-policy max-age 90
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure account password-policy min-age

```
configure account [all | name] password-policy min-age [num_days | none]
```

Description

Configures a minimum password lifespan.

Syntax Description

all	Applies the configuration to all accounts.
<i>name</i>	Applies the configuration to the account of the specified name.
password-policy	Specifies configuring the account password policy.
min-age	Specifies the minimum lifespan of passwords.
<i>num_days</i>	Specifies a minimum lifespan of passwords in days. Range is 1 to 365 days.
none	Specifies no limitation on the minimum lifespan for passwords (default).

Default

The default is no minimum lifespan for passwords.

Usage Guidelines

Similar to the `configure account [all | name] password-policy max-age [num_days | none]` command, which requires a password change after a configurable number of days, this command configures a lifespan, or minimum age. Having a minimum lifespan ensures that multiple password changes are not made in quick succession.

In the case of conflicting settings between these two commands, a setting requiring a password change overrides a setting that prohibits a password change. For example, if **max-age** is set to 10 days, thus requiring a password change in 10 days, and a **min-age** is set to 20 days, attempting to forbid a password change until 20 days, the configuration to change the password after 10 days takes precedence over the configuration to not change the password for 20 days.

To view the current selection for the minimum lifespan for passwords, use the `show accounts password-policy` command.

Example

The following example sets a minimum lifespan of 10 days for all accounts:

```
# configure account all password-policy min-age 10
```

History

This command was first available in ExtremeXOS 30.7.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure account password-policy min-different-characters

```
configure account [all | name] password-policy min-different-characters
    [count]
```

Description

When changing a password, configures the number of characters in the revised password that must be changed from the existing password.

Syntax Description

all	Applies the configuration to all accounts.
<i>name</i>	Applies the configuration to the account of the specified name.
password-policy	Specifies configuring the account password policy.
min-different-characters	Specifies the minimum number of different characters between a previous and new password.
<i>count</i>	Specifies the number of characters required to be different between the previous and new password. The range is 0-16. The default is 0.

Default

The default for the minimum number of different characters is 0.

Usage Guidelines

This command allows you to configure a variable number of characters that must be changed from the existing password when a new password is created. If the new password is longer than the original password, the “extended” characters of the new password are counted as different from the prior password. If the new password is shorter than the existing password, only the new password characters determine the number of characters that are different.

For example, if you specify 8 for the number of characters that must be different:

This combination would count as 9 different characters, and would pass:

- Current password: MyChoice
- New password: MyChoiceButLonger

This combination would count as zero different characters, and would fail:

- Current Password: MyChoiceButLonger
- New password: MyChoice

To view the current selection for the minimum number of different characters for changed passwords, use the `show accounts password-policy` command.

Example

The following example configures the minimum number of different characters for changed passwords to be "3" for all accounts:

```
# configure account all password-policy min-different-characters 3
```

History

This command was first available in ExtremeXOS 30.7.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure account password-policy min-length

```
configure account [all | name] password-policy min-length
                 [num_characters | none]
```

Description

Requires a minimum number of characters for passwords.

Syntax Description

all	Specifies all accounts (and future users).
<i>name</i>	Specifies an account name.

<i>num_characters</i>	Specifies the minimum number of characters required for the password. The range is 1–32 characters. Note: If you configure the configure account password-policy char-validation parameter, the minimum length is eight characters.
none	Resets password to accept a minimum of 0 characters. Note: If you configure the configure account encrypted parameter, the minimum length is eight characters.

Default

N/A.

Usage Guidelines

Use this command to configure a minimum length restriction for all passwords for specified accounts.

This command affects the minimum allowed length for the next password; the current password is unaffected.

The minimum password length is configurable from 1–32 characters. Using the none option disables the requirement of minimum password length and returns the system to the default state (password minimum is 0 by default).



Note

If the account is configured to require a specific password format, the minimum is 8 characters. See [configure account password-policy char-validation](#) for more information.

Example

The following command requires a minimum of 8 letters for the password for the account management:

```
configure account management password-policy min-length 8
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure account privilege

```
configure account [all | name] privilege [admin | user]
```

Description

Changes the privileges of an existing user account.

Syntax Description

account	Login account.
all	Specifies all accounts.
<i>name</i>	Specifies a specific user account.
privilege	Change the account privilege.
admin	Administrative privilege.
user	User (non-administrative) privilege.

Default

None.

Usage Guidelines

If an account is changed, any sessions that are currently logged in with that account are cleared, and therefore forced to login again with the new privilege. If the specified account is logged in to a session that cannot be cleared, an error message appears. If the account privilege is not changed by the option selected in the command for the specified account(s) (account already has that privilege), the request is ignored and any sessions logged in with the account are not cleared.

If you attempt to remove administrative privileges from the sole account having administrative privilege, you receive an error message.

Example

The following example adds administrative privilege to an account called "my_name":

```
configure account my_name privilege admin
```

History

This command was first available in ExtremeXOS 22.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure auto-peering oneconfig bootprelay

```
configure auto-peering one-config bootprelay [add | delete] [ip_address  
| ipv6_address] vr vrname
```

Description

For Auto-peering, adds dynamic BOOTP relay servers that the DHCP relay agent uses to forward DHCP traffic received from host attachments.

Syntax Description

auto-peering	Specifies Auto-peering.
one-config	Specifies changing the BGP Auto-peering OneConfig configuration.
bootprelay	Specifies configuring dynamic BOOTP relay service applied to dynamically created VRFs.
add	Adds dynamic BOOTP relay server.
delete	Deletes dynamic BOOTP relay server.
<i>ip_address</i>	Specifies the IPv4 address of the BOOTP relay server.
<i>ipv6_address</i>	Specifies the IPv6 address of the BOOTP relay server.
vr	Specifies selecting the virtual router (VR) If you do not specify a VR or VRF, the current VR context is used.
<i>vrname</i>	Specifies the VR name.

Default

If you do not specify a VR or VRF, the current VR context is used.

Usage Guidelines

Dynamic BOOTP relay services are configured after the dynamic VRF is installed on the device.

Example

The following example adds dynamic BOOTP relay server at "50.1.101.105" for VR "red":

```
# configure auto-peering oneconfig bootprelay add 50.1.101.105 vr red
```

History

This command was first available in ExtremeXOS 31.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure auto-peering oneconfig id

```
configure auto-peering one-config id [none | id]
```

Description

Configures the ID used by each device when automatically forming an adjacency with an BGP Auto-peering neighbor.

Syntax Description

auto-peering	Specifies Auto-peering.
one-config	Specifies changing the BGP Auto-peering OneConfig configuration.
id	Specifies configuring the BPG Auto-peering ID that devices advertise to neighbors.
none	Removes the current ID.
<i>id</i>	Specifies the ID.

Default

N/A.

Usage Guidelines

All devices in an Auto-peering cluster have the same ID. Neighbors with different IDs must match the Remote ID table.

Example

The following example configures the ID as "123":

```
# configure auto-peering one-config id 123
```

History

This command was first available in ExtremeXOS 31.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure auto-peering one-config iproute

```
configure auto-peering one-config iproute [add | delete] [host | hostv6]  
    [[ipaddress {netmask} | ipNetmask] gateway | ipNetmaskv6 gatewayv6]  
    {vr vrname }
```

Description

Configures a list of OneConfig dynamic Auto-peering static routes.

Syntax Description

auto-peering	Specifies Auto-peering.
one-config	Specifies changing the BGP Auto-peering OneConfig configuration.
iproute	Specifies configuring the dynamic IP route configured when the host IP address attaches to the device.
add	Adds route.
delete	Deletes routes.
host	Specifies the IPv4 host address.
hostv6	Specifies the IPv6 host address.
<i>ipaddress</i>	Specifies the IPv4 address.
<i>netmask</i>	Specifies the IPv4 netmask.
<i>ipNetmask</i>	Specifies the IPv4 netmask/mask.
<i>gateway</i>	Specifies the IPv4 gateway address.
<i>ipNetmaskv6</i>	Specifies the IPv6 address/netmask.
<i>\gatewayv6</i>	Specifies the IPv6 gateway address.
vr	Specifies selecting the virtual router (VR) If you do not specify a VR or VRF, the current VR context is used.
<i>vrname</i>	Specifies the VR name.

Default

If you do not specify a VR or VRF, the current VR context is used.

Usage Guidelines

The switch holds the static route in the dynamic database until the host attaches to the switch. The routes are applied in the route table when the host IP address is discovered on a local access port. They are removed when the host is no longer connected to the device. To see if the host routes are properly installed, use the `show iproute {ipv4} {priority | vlan vlan_name | permanent | ip_address netmask | summary} {multicast | unicast} {vr vrname}}` command.

Example

The following example adds the static IP route for host 50.1.102.101, IP/netmask 60.1.1.0/24, gateway 50.1.102.101 on VR "red":

```
# configure auto-peering oneconfig iproute add 50.1.102.101 60.1.1.0/24 50.1.102.101 vr
red
```

History

This command was first available in ExtremeXOS 31.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure auto-peering one-config nsi-id

```
configure auto-peering one-config nsi-id id type [nsi | vrf] [add
| delete] [[ipaddress {netmask} | ipNetmask ] | ipNetmaskv6] {vr
vrname }
```

Description

Specifies a list of OneConfig dynamic auto-peering services.

Syntax Description

auto-peering	Specifies Auto-peering.
one-config	Specifies changing the BGP Auto-peering OneConfig configuration.
nsi-id	Specifies setting the Network Service (NSI) ID for dynamic L3 configuration
<i>id</i>	Specifies the NSI ID.
type	Specifies configuring the NSI ID type.
nsi	Specifies NSI ID type as NSI.
vrf	Specifies NSI ID type as VRF.
add	Adds subnet.
delete	Deletes subnet.
<i>ipaddress</i>	Specifies the IP address.
<i>netmask</i>	Specifies the IP netmask.
<i>ipNetmask</i>	Specifies the IPv4 address/netmask.
<i>ipNetmaskv6</i>	Specifies IPv6 address/netmask
vr	Specifies selecting the virtual router (VR) If you do not specify a VR or VRF, the current VR context is used.
<i>vrname</i>	Specifies the VR name.

Default

If you do not specify a VR or VRF, the current VR context is used.

Usage Guidelines

The configuration is held on each node and is inactive until a VLAN/NSI binding is made to a dynamic VLAN. After the NSI is discovered, the configuration associated with the NSI on the VLAN is applied. This consists of creating the VRF as needed, applying IPv4 and IPv6 prefixes, enabling Anycast IP, and enabling IP Forwarding on the mapped VLAN. The mechanism to assign VLAN to NSI mapping is done

by following methods: policy by RADIS, and Fabric Attach; the VXLAN virtual network and VNI/NSI is dynamically created by the VXLAN service.

Example

The following example configures an NSI ID "1000" for type "VRF" on VR "red":

```
# configure auto-peering one-config nsi 1000 type VRF add vr red
```

History

This command was first available in ExtremeXOS 31.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure auto-peering oneconfig overlay

```
configure auto-peering one-config overlay [add | delete] server [address  
| addressv6] {type bgp-rr } {id id} {password [none | {encrypted}  
password] }
```

Description

Configures the Auto-peering overlay service database entries.

Syntax Description

auto-peering	Specifies Auto-peering.
one-config	Specifies changing the BGP Auto-peering OneConfig configuration.
overlay	Specifies overlay database containing centralized EVPN information.
add	Overlay add server.
delete	Overlay delete server.
server	Specifies the server.
<i>address</i>	Specifies the IPv4 server address/
<i>addressv6</i>	Specifies the IPv6 server address/
type	Specifies the overlay database type.
bgp-rr	Specifies overlay database type BGP route reflector.
id	Specifies the configuring the overlay database ID.
<i>id</i>	Specifies the overlay database ID. <ul style="list-style-type: none"> For BGP-RR—AS number. For Redis—port number.
password	Specifies configuring the password for the overlay database.

none	Removes the current password.
encrypted	Specifies encrypted format for the password.
<i>password</i>	Specifies the password/secret key.

Default

N/A.

Usage Guidelines

The overlay database allows the VXLAN edge technologies to dynamically span across brown field networks. For BGP-RR the ID represents the AS number, and for Redis, it represents the port to connect to.

Example

The following example configures the overlay database of type BGP-RR at "50.1.133.105" with ID "2000":

```
# configure auto-peering oneconfig overlay add server 50.1.133.105 type BGP-RR id 2000
```

History

This command was first available in ExtremeXOS 31.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure auto-peering one-config password

```
configure auto-peering one-config password [none | {encrypted}  
      tcpPassword]
```

Description

Configures the auto-peering TCP MD5 password devices will configure in the TCP MD5.

Syntax Description

auto-peering	Specifies Auto-peering.
one-config	Specifies changing the BGP Auto-peering OneConfig configuration.
password	Configures the password for Auto peering neighbors.
none	Removes the current password
encrypted	Specifies that the password is in encrypted format.
<i>tcpPassword</i>	Specifes the TCP MD5 password/secret-key.

Default

N/A.

Usage Guidelines

Changing auto peering password might cause peers to be disconnected if passwords do not match.

Example

The following example sets the password to "123":

```
# configure auto-peering one-config password 123
```

History

This command was first available in ExtremeXOS 31.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure auto-peering one-config remote id

```
configure auto-peering one-config remote id [add | delete] id {password  
[none | {encrypted} tcpPassword] }
```

Description

Configures a list of unique values that identify the remote Auto-peering devices to which this device can also automatically form an adjacency.

Syntax Description

auto-peering	Specifies Auto-peering.
one-config	Specifies changing the BGP Auto-peering OneConfig configuration.
remote	Specifies a remote device connecting in the Auto-peering network.
id	Specifies configuring the BPG Auto-peering ID that devices advertise to neighbors.
add	Specifies adding a remote ID.
delete	Specifies deleting a remote ID.
<i>id</i>	Specifies the ID.
password	Configures the password for Auto peering neighbors.
none	Removes the current password
encrypted	Specifies that the password is in encrypted format.
<i>tcpPassword</i>	Specifies the TCP MD5 password/secret-key.

Default

N/A.

Usage Guidelines

Not specifying a password will result in unsecure peering. Remote ID neighbors can use MD5 passwords for neighboring security. The lower ID password is used.

Example

The following example adds the remote Auto-peering device with ID "2222" and sets the password as "123":

```
# configure auto-peering one-config remote id add 2222 password 123
```

History

This command was first available in ExtremeXOS 31.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure automation edge connect/disconnect

```
configure automation edge [connect | disconnect] database database_name
```

Description

Connects or disconnects an Automation Edge remote VXLAN (Virtual Extensible LAN) network identifier (VNI)-device database.

Syntax Description

automation	Designates configuring Automation Edge VXLAN VNI-device database information.
edge	Designates configuring Automation Edge VXLAN VNI-device database information.
connect	Designates connecting to the specified database.
disconnect	Designates disconnecting from the specified database.
database	Designates connecting or disconnecting a database.
<i>database_name</i>	Sets the database name that you are connecting or disconnecting.

Default

N/A.

Usage Guidelines

To connect to another new database, an existing database must be disconnected first. To view the connected database, use the command `show database database_name`

Example

The following example connects the database "database1":

```
# configure automation edge connect database database1
```

The following example disconnects the database "database1":

```
# configure automation edge disconnect database database1
```

History

This command was first available in ExtremeXOS 31.1 as a demonstration feature.

This command is fully supported in ExtremeXOS 31.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure banner

```
configure banner {after-login | { before-login } { acknowledge } |  
                 before-login {acknowledge} save-to-configuration}
```

Description

Configures the banner string to be displayed for CLI screens.

Syntax Description

after-login	Specifies that a banner be displayed after login.
before-login	Specifies that a banner be displayed before login.
acknowledge	Require acknowledgement of the banner before login.
save-to-configuration	Save the before login banner to the configuration file as well as non-volatile memory.

Default

N/A.

Usage Guidelines

Use this command to configure two types of banners:

- A banner for a CLI session that displays before login.
- A banner for a CLI session that displays after login.

If no optional parameters are specified, the command defaults to configuring a banner that is displayed before the CLI session login prompt.

For each CLI session banner, you can enter up to 24 rows of 79-column text.

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.



Note

The system does not wait for a keypress when you use SSH for access; this only applies to the serial console login sessions and Telnet sessions.

To disable the acknowledgement feature, use the `configure banner` command omitting the `acknowledge` parameter.

To display any configured banners, use the `show banner` command.

To unconfigure one or more configured banners, use the `unconfigure banner` command.

Example

The following example add the text "test" before the pre-login prompt:

```
# configure banner before-login
test

# logout
Do you wish to save your configuration changes to primary.cfg? (y/N)
Y
test

login:
# show banner

Before-Login banner:
test

Acknowledge: Disabled
Save to      : Non-volatile memory only

After-Login banner:
```

History

This command was first available in ExtremeXOS 10.1.

The `acknowledge` parameter was added in ExtremeXOS 11.5.

The `after-login` option was added in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure bfd vlan

```
configure bfd vlan vlan_name [{detection-multiplier multiplier}  
  {receive-interval rx_interval} {transmit-interval tx_interval}]
```

Description

Configures BFD transmit (TX) and receive (RX) intervals and multipliers on the VLAN.

Syntax Description

<i>vlan_name</i>	Specifies the VLAN.
<i>multiplier</i>	Specifies the detection multiplier. The range is 1 to 255.
<i>rx_interval</i>	Specifies the receive interval for control packets in milliseconds. The range is 100 to 4294967 ms. (3 to 4294967 ms if hardware assist is enabled).
<i>tx_interval</i>	Specifies the transmit interval for control packets in milliseconds. The range is 100 to 4294967 ms. (3 to 4294967 ms if hardware assist is enabled).

Default

The default value for RX and TX intervals is 1000 ms.

The default value for the detection-multiplier is 3.

Usage Guidelines

Use this command to configure BFD.

Use the `show bfd vlan` command to display the current settings.

Example

The following command configures a transmit and receive interval of 2000 ms and a detection multiplier of 2 on the VLAN `vlan1`:

```
# configure bfd vlan vlan1 detection-multiplier 2 receive-interval 2000 transmit-interval 2000
```

Receive interval of 0

An *rx_interval* value of 0 means that this system does not want to receive any periodic BFD Control packets. A system may transmit a value of 0 for the Required MinRX Interval to indicate that the remote system should send no packets.

```
# configure bfd vlan vlan1 detection-multiplier 2 receive-interval 0 transmit-interval 2000
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure bfd vlan authentication

```
configure bfd vlan vlan_name authentication [none | simple-password  
{encrypted encrypted_password | password }]
```

Description

Configures authentication for BFD on a VLAN.

Syntax Description

<i>vlan_name</i>	Specifies the VLAN name.
none	Specifies that no authentication is to be used. (Default)
encrypted	Indicates that the password is already encrypted.
<i>password</i>	Specifies a simple password to use to authenticate.

Default

The authentication default is none.

Usage Guidelines

Use this command to configure authentication for BFD on a VLAN using a password or specify that none is required.

Use the `show bfd vlan` command to display the authentication setting.

The encrypted keyword is primarily for the output of the show configuration command, so that the password is not revealed in the command output. Do not use it to set the password

Example

The following command configures authentication using the password password:

```
# configure bfd vlan vlan1 authentication simple-password password
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure bgp add aggregate-address

```
configure bgp add aggregate-address {address-family [ipv4-unicast |  
    ipv4-multicast | ipv6-unicast | ipv6-multicast]} ipaddress/masklength  
    {as-match | as-set} {summary-only} {advertise-policy policy}  
    {attribute-policy policy}
```

Description

Configures a *BGP* aggregate route.

Syntax Description

address-family	Specifies an IPv4 or IPv6 unicast or multicast address family.
<i>ipaddress/masklength</i>	Specifies an IP network address and mask length.
as-match	Generates autonomous system sequence path information (order of AS numbers in AS_PATH is preserved).
as-set	Generates autonomous system set path information (order of AS numbers in AS_PATH is not preserved).
summary-only	Specifies to send only aggregated routes to the neighbors.
advertise-policy	Specifies the policy used to select routes for this aggregated route.
attribute-policy	Specifies the policy used to set the attributes of the aggregated route.

Default

If no address family is specified, IPv4 unicast is the 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.

Before you can create an aggregate route, you must enable BGP aggregation using the following command:

```
enable bgp aggregation
```

If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.



Note

If the specified address is an IPv4 address, an IPv4 address family must be specified with the command. If the specified address is an IPv6 address, an IPv6 address family must be specified with the command.

BGP supports overlapping routes. For example, you can configure both of the following aggregate addresses:

- 192.0.0.0/8
- 192.168.0.0/16

After you create an aggregate route, the aggregate route remains inactive until BGP receives a route with an IP address and mask that conforms to an aggregate route. When a conforming route is received, the aggregate route becomes active and is advertised to BGP neighbors. If the summary-only option is specified, only the aggregate route becomes active and is advertised. If the summary-only option is omitted, any conforming aggregate routes and the received route are advertised to BGP neighbors.

Example

The following command configures a BGP aggregate route:

```
configure bgp add aggregate-address 192.1.1.4/30
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for overlapping aggregate addresses was added in ExtremeXOS 12.1.3.

Support for IPv6 was added in ExtremeXOS 12.6-BGP.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure bgp add confederation-peer sub-AS-number

```
configure bgp add confederation-peer sub-AS-number number
```

Description

Adds a sub-AS to a confederation.

Syntax Description

<i>number</i>	Specifies a sub-AS number of the confederation. The range is 1 to 4294967295.
---------------	---

Default

N/A.

Usage Guidelines

Before you can add a sub-AS to a confederation on the switch, you must disable any *BGP* neighbor sessions that are configured with the same AS number as a remote AS number. To disable BGP neighbor sessions, use the following command:

```
disable bgp neighbor [remoteaddr | all]
```

Invoke the `configure bgp add confederation-peer sub-AS-number` command multiple times to add multiple sub-ASs.

IBGP 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, all BGP speakers in each sub-AS must be fully-meshed. The confederation is advertised to other networks as a single AS.

The AS number is a 4-byte AS number in either the ASPLAIN or the ASDOT format as described in RFC 5396, Textual Representation of Autonomous System (AS) Numbers.

Example

The following example adds one sub-AS to a confederation using the ASPLAIN 4-byte AS number format:

```
configure bgp add confederation-peer sub-AS-number 65536
```

The following example adds one sub-AS to a confederation using the ASDOT 4-byte AS number format:

```
configure bgp add confederation-peer sub-AS-number 1.15
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for 4-byte AS numbers was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure bgp add network

```
configure bgp add network {address-family [ipv4-unicast | ipv4-multicast
| ipv6-unicast | ipv6-multicast]} ipaddress/masklength {network-policy
policy}
```

Description

Adds a network to be originated from this router.

Syntax Description

address-family	Specifies an IPv4 or IPv6 unicast or multicast address family.
<i>ipaddress/masklength</i>	Specifies an IP network address and mask length.
<i>policy</i>	Name of policy to be associated with network export. Policy can filter and/or change the route parameters.

Default

If no address family is specified, IPv4 unicast is the default.

N/A.

Usage Guidelines

The network must be present in the routing table.

Using the export command to redistribute routes complements the redistribution of routes using the `configure bgp add network` command. The `configure 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.

If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.

**Note**

If the specified address is an IPv4 address, an IPv4 address family must be specified with the command. If the specified address is an IPv6 address, an IPv6 address family must be specified with the command.

Example

The following command adds a network to be originated from this router:

```
configure bgp add network 192.1.1.16/32
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure bgp as-display-format

```
configure bgp as-display-format [asdot | asplain]
```

Description

Configures the AS number format displayed in show commands.

Syntax Description

asdot	Specifies the ASDOT format.
asplain	Specifies the ASPLAIN format.

Default

N/A.

Usage Guidelines

The ASPLAIN and ASDOT formats are described in RFC 5396, Textual Representation of Autonomous System (AS) Numbers.

Example

The following command selects the ASDOT 4-byte AS number format:

```
configure bgp as-display-format asdot
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the *BGP* feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure bgp as-number

```
configure bgp AS-number number
```

Description

Changes the local AS number used by *BGP*.

Syntax Description

<i>number</i>	Specifies a local AS number. The range is 1 to 4294967295.
---------------	--

Default

N/A.

Usage Guidelines

BGP must be disabled before the AS number can be changed.

This command applies to the current VR or VRF context.

The AS number is a 4-byte AS number in either the ASPLAIN or the ASDOT format as described in RFC 5396, Textual Representation of Autonomous System (AS) Numbers.

Example

The following command specifies a local AS number using the ASPLAIN 4-byte AS number format:

```
configure bgp AS-number 65551
```

The following command specifies a local AS number using the ASDOT 4-byte AS number format:

```
configure bgp AS-number 1.15
```



Note

To remove the configured bgp as-number, assign as-number value as 0, i.e. configure bgp AS-number 0.

The following command configures the BGP router ID:

```
configure bgp routerid
```



Note

To remove the configured bgp routerid, give routerid value as 0.0.0.0 i.e. configure bgp routerid 0.0.0.0.

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for 4-byte AS numbers was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure bgp cluster-id

```
configure bgp cluster-id cluster-id
```

Description

Configures the local cluster ID.

Syntax Description

<i>cluster-id</i>	Specifies a 4 byte field used by a route reflector to recognize updates from other route reflectors in the same cluster. The range is 0 - 4294967295.
-------------------	---

Default

N/A.

Usage Guidelines

BGP must be disabled before the cluster ID can be changed.

Used when multiple route reflectors are used within the same cluster of clients.

Example

The following command appends a BGP route reflector cluster ID to the cluster list of a route:

```
configure bgp cluster-id 40000
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure bgp confederation-id

```
configure bgp confederation-id confederation-id
```

Description

Specifies a *BGP* routing confederation ID.

Syntax Description

<i>confederation-id</i>	Specifies a routing confederation identifier, which is a 4-byte AS number in the range of 1 to 4,294,967,295.
-------------------------	---

Default

N/A.

Usage Guidelines

IBGP requires that networks 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.

The confederation ID is a 4-byte AS number in either the ASPLAIN or the ASDOT format as described in RFC 5396, Textual Representation of Autonomous System (AS) Numbers.

BGP must be disabled before the confederation ID can be changed.

Use a confederation ID of 0 to indicate no confederation. You cannot unconfigure the confederation ID while confederation peers are configured. You must delete the confederation peers before you unconfigure the confederation ID.

Example

The following command specifies a BGP routing confederation ID using the ASPLAIN 4-byte AS number format:

```
configure bgp confederation-id 65551
```

The following command specifies a BGP routing confederation ID using the ASDOT 4-byte AS number format:

```
configure bgp confederation-id 1.15
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for 4-byte AS numbers was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure bgp delete aggregate-address

```
configure bgp delete aggregate-address {address-family [ipv4-unicast  
| ipv4-multicast | ipv6-unicast | ipv6-multicast]} [ ipaddress/  
masklength | all]
```

Description

Deletes one or all *BGP* aggregated routes.

Syntax Description

address-family	Specifies an IPv4 or IPv6 unicast or multicast address family.
<i>ipaddress/masklength</i>	Specifies an IP network address and netmask length.
all	Specifies all aggregated routes in the specified address family. If you do not specify an address family, all aggregated routes in all address families are deleted.

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.

If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.



Note

If the specified address is an IPv4 address, an IPv4 address family must be specified with the command. If the specified address is an IPv6 address, an IPv6 address family must be specified with the command.

Example

The following command deletes a BGP aggregate route:

```
configure bgp delete aggregate-address 192.1.1.4/30
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure bgp delete confederation-peer sub-AS-number

```
configure bgp delete confederation-peer sub-AS-number number
```

Description

Specifies a sub-AS that should be deleted from a confederation.

Syntax Description

sub-AS-number	Specifies a sub-AS.
----------------------	---------------------

Default

N/A.

Usage Guidelines

Before you can change the configuration with this command, you must disable the [BGP](#) neighbors in the confederation using the following command:

```
disable bgp neighbor [remoteaddr | all]
```

Example

The following command deletes a sub-AS from a confederation using the ASPLAIN 4-byte AS number format:

```
configure bgp delete confederation-peer sub-AS-number 65551
```

The following command deletes a sub-AS from a confederation using the ASDOT 4-byte AS number format:

```
configure bgp delete confederation-peer sub-AS-number 1.15
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for 4-byte AS numbers was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure bgp delete network

```
configure bgp delete network {address-family [ipv4-unicast | ipv4-multicast | ipv6-unicast | ipv6-multicast]} [all | ipaddress/masklength]
```

Description

Deletes a network to be originated from this router.

Syntax Description

address-family	Specifies an IPv4 or IPv6 unicast or multicast address family.
all	Specifies all networks for the specified address family. If no address family is specified, all networks for all address families are deleted.
<i>ipaddress/masklength</i>	Specifies an IP network address and netmask length.

Default

N/A.

Usage Guidelines

If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.



Note

If the specified address is an IPv4 address, an IPv4 address family must be specified with the command. If the specified address is an IPv6 address, an IPv6 address family must be specified with the command.

Example

The following command deletes a network to be originated from this router:

```
configure bgp delete network 192.1.1.12/30
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 [BGP](#).

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure bgp evpn ignore-as

```
configure bgp evpn ignore-as [on | off]
```

Description

Configures treatment of the anonymous system (AS) part of the route target.

Syntax Description

bgp	BGP capability.
evpn	EVPN protocol.
ignore-as	Configure treatment of the AS part of the route target.
on	Specifies that route target matching ignore the AS part of the route target (default).
off	Specifies that route target matching not ignore the AS part of the route target.

Default

By default, route target matching ignores the AS part of the route target.

Usage Guidelines

To view the current setting for ignore-as, use the `show bgp evpn` command.

Example

The following example configures route target matching to ignore the AS as part of the route target:

```
# configure bgp evpn ignore-as on
```

History

This command was first available in ExtremeXOS 31.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure bgp evpn instance rd

```
configure bgp evpn instance evpn_instance_name rd [rd_value | auto]
```

Description

Configures route distinguishers for an EVPN instance.

Syntax Description

bgp	BGP capability.
evpn	EVPN protocol.
instance	Specifies configuring an EVPN instance
<i>evpn_instance_name</i>	Specifies name of the EVPN instance.
rd	Specifies configuring route distinguisher.
<i>rd_value</i>	Route distinguisher in format <admin>:<assigned number>.
auto	Specifies auto-derived route distinguisher values (default).

Default

By default, auto-derived route distinguisher values are used.

Example

The specifies auto-derived route distinguisher values for the EVPN instance "my_evpn":

```
# configure bgp evpn instance my_evpn rd auto
Warning: Changing RD value for EVPN instance my_evpn from  to 'auto calculated' instance
will be reset
```

History

This command was first available in ExtremeXOS 30.2.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure bgp evpn instance route-target

```
configure bgp evpn instance evpn_instance_name route-target {import |  
export | both} [add | delete] route_target
```

Description

Configures route targets for an EVPN instance.

Syntax Description

bgp	BGP capability.
evpn	EVPN protocol.
instance	Specifies configuring an EVPN instance
<i>evpn_instance_name</i>	Specifies name of the EVPN instance.
route-target	Designates setting the route target association. Default is autoderived import and export.
import	Selects import route target.
export	Selects export route target.
both	Specifies import and export route target mode (default).
add	Adds a route target.
delete	Deletes a route target.
<i>route_target</i>	Route target in format <global-admin-value>:<local-admin-value>.

Default

By default, if you do not specify route target, then the auto-derived values are used.

If you do not specify, import and export route target mode applies

Usage Guidelines

For EBGp applications of EVPN, the auto-derived values will not match between BGP peers since local autonomous system (AS) is used in the derivation, and these differ between EBGp peers.

Note that the route target mode (import, export, or both) is automatically adjusted depending on configuration. For example, if an “import” target exists and you add an “export” target for the same value, the mode is automatically changed to “both”. Similarly, an entry can be deleted by mode. For example, if an entry has mode of “both” and you delete the “import” target of the same value, the entry is not deleted, instead its mode is changed to “export”. An attempt to delete an entry that does not exist (value or mode) produces an error message and no action is taken. For example, if you attempt to delete a route target using “both”, but the configured entry was only configured as “import” an error message appears, and no action is taken.

Example

The following example configures for instance "my_evpn" route target both mode:

```
# configure bgp evpn instance my_evpn route-target both
```

History

This command was first available in ExtremeXOS 30.2.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure bgp evpn instance vxlan

```
configure bgp evpn instance evpn_instance_name vxlan vni [vni_value |  
none]
```

Description

Adds or deletes a virtual extensible local area network (VXLAN) virtual network identifier (VNI) to an EVPN instance.

Syntax Description

bgp	BGP capability.
evpn	EVPN protocol.
instance	Specifies configuring an EVPN instance
<i>evpn_instance_name</i>	Specifies name of the EVPN instance.
vxlan	Specifies termination.
vni	Specifies adding a VXLAN VNI to an EVPN instance.
<i>vni_name</i>	Specifies the VNI (range = 1-16,777,215).
none	Removes existing VNI setting for this EVPN instance.

Default

N/A.

Example

The following example adds the VXLAN VNI "12345" to an EVPN instance named "my_evpn":

```
# configure bgp evpn instance my_evpn vxlan vni 12345
```

The following example removes the existing VXLAN VNI associated with the EVPN instance named "my_evpn":

```
# configure bgp evpn instance my_evpn vxlan vni none
```

History

This command was first available in ExtremeXOS 30.2.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure bgp evpn l3vni

```
configure bgp evpn l3vni [vni_value | none] vr vr_name
```

Description

Configures a Layer 3 virtual network identifier (VNI) and binds it to a VPN virtual router interface (VRF).

Syntax Description

bgp	Specifies changing the BGP configuration.
evpn	Specifies changing the EVPN configuration.
l3vni	Specifies changing the integrated routing and bridging IP VRF VNI configuration.
<i>vni_value</i>	Sets the VNI to a value between 1 and 16,777,215.
none	Removes an existing VNI setting for this EVPN L3VNI.
vr	Specifies binding the VNI to a VRF.
<i>vr_name</i>	Sets the VRF name.

Default

N/A.

Example

The following example binds VNI 100 to VRF "vrf1":

```
# configure bgp evpn l3vni 100 vr vrf1
```

History

This command was first available in ExtremeXOS 30.6.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure bgp export shutdown-priority

```
configure bgp export route_type {{address-family} address_family}
shutdown-priority number
```

Description

Configures the shutdown priority for IGP export.

Syntax Description

<i>route_type</i>	Specifies the <i>BGP</i> export route type.
address-family	Specifies an IPv4 or IPv6 unicast or multicast address family.
<i>number</i>	Specifies the shutdown priority. The range is 0 - 65,535.

Default

The default value is 2048.

If no address family is specified, IPv4 unicast is the default.



Note

You must specify an IPv6 address family for an IPv6 peer, because an IPv6 peer does not support the default IPv4 unicast address family. Similarly, if you specify an IPv4 peer and an address family in the command, an IPv4 address family must be specified.

Usage Guidelines

To export IPv6 protocols to BGP, you must specify an IPv6 address family.



Note

This command is not currently supported, and is not recommended for use.

Higher priority values lower the chance of an IGP export to be automatically disabled in case BGP or the system goes to a low memory condition.



Note

For this command to execute, the specified protocol must support the specified address family. For example, the command fails if you specify *OSPF* and the IPv6 unicast address family. You can specify blackhole, direct, static, and IS-IS routes with IPv4 or IPv6 address families.

Example

The following command configures the shutdown priority of BGP exported OSPF routes to 1000:

```
configure bgp export ospf shutdown-priority 1000
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure bgp import-policy

```
configure bgp import-policy [policy-name | none]
```

Description

Configures the import policy for *BGP*.

Syntax Description

<i>policy-name</i>	Specifies the policy.
none	Specifies no policy.

Default

N/A.

Usage Guidelines

Use the none keyword to remove a BGP import policy.

An import policy is used to modify route attributes while adding BGP routes to the IP route table.

Example

The following command configures a policy imprt_plcy for BGP:

```
configure bgp import-policy imprt_plcy
```

The following command unconfigures the import policy for BGP:

```
configure bgp import-policy none
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure bgp local-preference

```
configure bgp local-preference number
```

Description

Changes the default local preference attribute.

Syntax Description

<i>number</i>	Specifies a value used to advertise this router's degree of preference to other routers within the AS. Range is 0 to 2147483647.
---------------	--

Default

100.

Usage Guidelines

BGP must be disabled before the local preference attribute can be changed.

BGP selects routes based on the following precedence (from highest to lowest):

- higher weight
- higher local preference
- shortest length (shortest AS path)
- lowest origin code
- lowest MED
- route from external peer
- lowest cost to Next Hop
- lowest routerID

Local preference is used to determine a preferred exit point from an AS. Local preferences are exchanged throughout the AS. A change in the local-preference can result in a change in routing and forwarding of traffic leaving the AS.

Example

The following command changes the default local preference attribute to 500:

```
configure bgp local-preference 500
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure bgp maximum-as-path-length

```
configure bgp maximum-as-path-length [max-as-path | none]
```

Description

This command adds support for filtering *BGP* updates based on a specified maximum autonomous system path (AS-path) length. This support is on a per BGP instance basis (not per neighbor).

Syntax Description

maximum-as-path-length	Specifies setting the AS path length.
<i>max-as-path</i>	Value specifying the AS path length. Range is 1 to 1,500.
none	Specifies no maximum AS path length.

Default

N/A

Usage Guidelines

It can be desirable to protect the router against BGP updates with excessively long AS-paths to ensure memory is not exhausted. Any BGP updates that exceed this user-defined limit are dropped. This setting does not affect existing routes.

Example

The following example sets the AS-path to 500.

```
configure bgp maximum-as-path-length 500
```

History

This command was first available in ExtremeXOS 21.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure bgp maximum-paths

```
configure bgp maximum-paths max-paths
```

Description

Enables or disables the BGP ECMP (Equal Cost Multi Paths) feature and specifies the maximum number of paths supported on the current VR.

Syntax Description

<i>max-paths</i>	Specifies the maximum number of paths. The range is 1 to 64. The value 1 disables BGP ECMP. A value greater than 1 enables BGP ECMP and specifies the maximum number of paths.
------------------	--

Default

One. BGP ECMP is disabled.

Usage Guidelines

This command triggers the BGP decision process, causing BGP to re-install the entire BGP routing table into the IP forwarding table. This activity requires a significant amount of switch processor resources, so we recommend that you enable or disable the BGP ECMP feature before enabling the BGP protocol globally on a VR. To ensure that BGP ECMP routes are programmed in the hardware, enter the enable iproute sharing command.



Note

BGP must be disabled before you can change the configuration with this command.

Example

The following command enables BGP ECMP and sets the maximum number of paths to 4 (the maximum number of possible paths is 64):

```
configure bgp maximum-paths 4
```

History

This command was first available in ExtremeXOS 12.1.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure bgp med

```
configure bgp med [none | bgp_med]
```

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

none	Specifies not to use a multi-exist-discriminator number.
<i>bgp_med</i>	Specifies a multi-exit-discriminator number. The range is 0-2147483647.

Default

N/A.

Usage Guidelines

BGP selects routes based on the following precedence (from highest to lowest):

- higher weight
- higher local preference
- shortest length (shortest AS path)
- lowest origin code
- lowest MED
- route from external peer

- lowest cost to Next Hop
- lowest routerID



Note

BGP must be disabled before you can change the configuration with this command.

Example

The following command configures the metric to be included in the MED path attribute:

```
configure bgp med 3
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure bgp neighbor allowas-in

```
configure bgp neighbor [all | remoteaddr] {address-family [ipv4-unicast
| ipv4-multicast | ipv6-unicast | ipv6-multicast | vpnv4] } allowas-in
{max-as-occurrence as-count}
```

Description

Configures EBGP to receive and accept a looped EBGP route from the specified neighbor, provided the number of occurrences of local AS number in AS-Path is less than or equal to the value of *as-count*.

Syntax Description

all	Specifies that the configuration applies to all neighbors in the specified address family. If no address family is specified or if an IPv4 address is specified, the configuration applies to all IPv4 neighbors. If an IPv6 address family is specified, the configuration applies to all IPv6 neighbors.
<i>remoteaddr</i>	Specifies the IPv4 or IPv6 address of a <i>BGP</i> neighbor. The switch uses the IP address format to determine if the address is an IPv4 or IPv6 address.
ipv4-unicast	Specifies the IPv4 unicast address family.
ipv4-multicast	Specifies an IPv4 multicast address family.

ipv6-unicast	Specifies the IPv6 unicast address family.
ipv6-multicast	Specifies an IPv6 multicast address family.
vpnv4	Specifies the VPNv4 address family for Layer 3 VPN support. This address family is applicable for PE to PE BGP neighbor sessions only. This keyword may prompt warning or error messages if executed for a regular BGP neighbor session or for a PE to CE neighbor session.
<i>as-count</i>	The maximum number of occurrences of local AS number in the received route AS-Path. If the number of occurrences of local AS number in AS-Path is more than as-count, the route is not accepted. The valid range is from 1-16.

Default

This feature is disabled by default.

If no as-count is specified, the as-count defaults to 3.

If no address family is specified and an IPv4 address is detected, IPv4 unicast is the default address family.

Usage Guidelines

In a hub and spoke configuration, it becomes necessary to accept an inbound EBGp route even though the route's AS-Path contains the receiver's own AS-number. In such network topologies, this feature can be enabled.



Note

A looped AS path is always allowed for IBGP, irrespective of the BGP configuration.

All EBGp routes with looped AS-Path are silently discarded by default.

This command applies to the current VR or VRF context.

If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.



Note

For an IPv6 peer, an IPv6 address family must be specified, because an IPv6 peer does not support IPv4 address families. If no address family is specified for an IPv6 peer, the default, the IPv4 unicast address family, applies and the command fails. Similarly an IPv4 peer only supports IPv4 address families and the command fails if an IPv6 address family is specified.

To configure this feature on Layer 3 VPNs, you must configure this feature in the context of the *MPLS*-enabled VR; this feature is not supported for BGP neighbors on the CE (VRF) side of the PE router.

Example

The following example enables BGP to accept looped BGP routes that contains a maximum of 6 occurrences of receiver's AS-number in AS-Path attribute:

```
configure bgp neighbor 192.162.17.54 allowas-in max-as-occurrence 6
```

History

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Support for Layer 3 VPNs was added in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure bgp neighbor alternate-local-as

```
configure bgp neighbor [all | remoteaddr] alternate-local-as asNumber
```

Description

Allows the local router to accept peering sessions intended for the specified alternate local autonomous system (AS).

Syntax Description

bgp	Specifies BGP.
neighbor	Specifies BGP neighbor.
all	Selects configuring all BGP neighbors.
<i>remoteaddr</i>	Selects configuring the specified BGP neighbor (IP address).
alternate-local-as	Allow alternate local AS number for peering with this neighbor
<i>asNumber</i>	AS number (0-4,294,967,295).

Default

N/A.

Usage Guidelines

This command provides configuration flexibility, particularly when peering with third-party devices that may use a different AS number than the ExtremeXOS device uses for auto-peering.

Example

The following example configures the BPG neighbor at 192.168.99.1 to use an alternate local AS "50":

```
# configure bgp neighbor 192.168.99.1 alternate-local-as 50
```

History

This command was first available in ExtremeXOS 30.2.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure bgp neighbor bfd

```
configure bgp {neighbor [all | remoteaddr ]} {bfd [on | off]}
```

Description

Enables or disables Bidirectional Forwarding Detection (BFD) protection of BGP peering sessions.

Syntax Description

all	Specifies that the configuration applies to all neighbors in the specified address family. If no address family is specified or if an IPv4 address is specified, the configuration applies to all IPv4 neighbors. If an IPv6 address family is specified, the configuration applies to all IPv6 neighbors.
<i>remoteaddr</i>	Specifies the IPv4 or IPv6 address of a BGP neighbor. The switch uses the IP address format to determine if the address is an IPv4 or IPv6 address.
bfd on off	Configures BFD detection for the specified neighbor(s).

Default

BFD is disabled on neighbor by default.

Usage Guidelines

You must disable a neighbor before configuring BFD.

Example

The following example enables BFD on neighbor 192.168.24.2:

```
# disable bgp neighbor 192.168.24.2
# configure bgp neighbor 192.168.24.2 bfd on
# enable bgp neighbor 192.168.24.2
```

History

This command was first available in ExtremeXOS 21.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure bgp neighbor connect-retry

```
configure bgp neighbor [all | remoteaddr] connect-retry seconds
```

Description

Configures the *BGP* neighbor retry timer.

Syntax Description

all	Specifies that the configuration applies to all neighbors in the specified address family. If no address family is specified or if an IPv4 address is specified, the configuration applies to all IPv4 neighbors. If an IPv6 address family is specified, the configuration applies to all IPv6 neighbors.
<i>remoteaddr</i>	Specifies the IPv4 or IPv6 address of a BGP neighbor. The switch uses the IP address format to determine if the address is an IPv4 or IPv6 address.
connect-retry	Specifies to configure the time interval between attempts to establish a TCP connection to a configured peer.
<i>seconds</i>	Specifies the retry time in seconds. Default is 30 seconds.

Default

Thirty (30) seconds when newly configured in ExtremeXOS 31.6 or later.

Usage Guidelines

Five specific values are supported: 10, 30, 60, 90, and 120 seconds.

Enter **show bgp neighbor detail** to display the current value of this command.

The neighbor retry timer was not configurable prior to ExtremeXOS 31.6 and was fixed at 120 seconds.

Example

The following command configures the BGP neighbor retry timer to 120 seconds:

```
configure bgp neighbor 192.168.1.22 connect-retry 120
```

History

This command was first available in ExtremeXOS 31.6.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure bgp neighbor dampening

```
configure bgp neighbor [all | remoteaddr] {address-family [ipv4-unicast | ipv4-multicast | ipv6-unicast | ipv6-multicast | vpn4]} dampening {{half-life half-life-minutes {reuse-limit reuse-limit-number suppress-limit suppress-limit-number max-suppress max-suppress-minutes} | policy-filter [policy-name | none]}}
```

Description

Configures the route flap dampening feature for a [BGP](#) neighbor.

Syntax Description

all	Specifies that the configuration applies to all neighbors in the specified address family. If no address family is specified or if an IPv4 address is specified, the configuration applies to all IPv4 neighbors. If an IPv6 address family is specified, the configuration applies to all IPv6 neighbors.
<i>remoteaddr</i>	Specifies the IPv4 or IPv6 address of a BGP neighbor. The switch uses the IP address format to determine if the address is an IPv4 or IPv6 address.
ipv4-unicast	Specifies the IPv4 unicast address family.
ipv4-multicast	Specifies an IPv4 multicast address family.
ipv6-unicast	Specifies the IPv6 unicast address family.
ipv6-multicast	Specifies an IPv6 multicast address family.
vpn4	Specifies the VPNv4 address family for Layer 3 VPN support. This address family is applicable for PE to PE BGP neighbor sessions only. Using this keyword may prompt warning or error messages if executed for a regular BGP neighbor session, or for a PE to CE neighbor session.
half-life	Specifies the dampening half life. Range is 1 to 45 minutes.

reuse-limit	Specifies the reuse limit. Range is 1 to 20000.
suppress-limit	Specifies the suppress limit. Range is 1 to 20000.
max-suppress	Specifies the maximum hold down time. Range is 1 to 255 minutes.
policy-filter	Specifies a policy.
none	Removes the configured policy.

Default

This feature is disabled by default.

If no address family is specified and an IPv4 address is detected, IPv4 unicast is the default address family.

Usage Guidelines

This command applies to the current VR or VRF context.

If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.



Note

For an IPv6 peer, an IPv6 address family must be specified, because an IPv6 peer does not support IPv4 address families. If no address family is specified for an IPv6 peer, the default, the IPv4 unicast address family, applies and the command fails. Similarly an IPv4 peer only supports IPv4 address families and the command fails if an IPv6 address family is specified.

The half life is the period of time, in minutes, during which the accumulated penalty of a route is reduced by half. The range is 1 to 45 minutes, and the default is 15 minutes.

The reuse limit is the penalty value below which a route is used again. The range is 1-20,000, and the default is 750.

The suppress limit is the penalty value above which a route is suppressed. The range is 1-20,000, and the default is 2,000.

The maximum hold down time is the maximum time a route can be suppressed, no matter how unstable it has been, as long as it no longer flaps. The range is 1-255 minutes, and the default is 4 * the half life.

If you change dampening parameters when routes are in suppressed or history state, the new dampening parameters apply only to routes in the active state. Routes in the suppressed or history state continue to use the old dampening parameters until they become active, at which time they use the updated dampening parameters.

Instead of explicitly configuring the dampening parameters using the command line, you can associate a policy using the policy-filter option. Multiple sets of parameters can be supplied using a policy.

Use the following command to disable route flap dampening for BGP neighbors:


```
configure bgp neighbor [remoteaddr | all] {address-family [ipv4-unicast
| ipv4-multicast | ipv6-unicast | ipv6-multicast | vpn4] } no-dampening
```

To configure this feature on Layer 3 VPNs, you must configure this feature in the context of the [MPLS](#)-enabled VR; this feature is not supported for BGP neighbors on the CE (VRF) side of the PE router.

Example

The following command configures route flap dampening to the BGP neighbor at 192.168.1.22 to the default values:

```
configure bgp neighbor 192.168.1.22 dampening
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Support for Layer 3 VPNs was added in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure bgp neighbor description

```
configure bgp neighbor [all | remoteaddr] description {description}
```

Description

Configures a description for a [BGP](#) neighbor.

Syntax Description

all	Specifies all IPv4 and IPv6 neighbors.
<i>remoteaddr</i>	Specifies the IPv4 or IPv6 address of a BGP neighbor.
<i>description</i>	Specifies a string used to describe the neighbor.

Default

The description is a NULL string by default.

Usage Guidelines

Use this command to attach a description to a BGP neighbor. This description is displayed in the output of the [show bgp neighbor](#) command when you specify the detail option, or when you specify a particular neighbor. Enclose the string in double quotes if there are any blank spaces in the string. The maximum length of the string is 56 characters.

If you do not specify the *description* parameter, the description is reset to the default.

This command applies to the current VR or VRF context.

Example

The following command configures the description for the BGP neighbor 192.168.1.22 to Toledo_5:

```
configure bgp neighbor 192.168.1.22 description Toledo_5
```

History

This command was first available in ExtremeXOS 11.0.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure bgp neighbor dont-allowas-in

```
configure bgp neighbor [all | remoteaddr] {address-family [ipv4-unicast | ipv4-multicast | ipv6-unicast | ipv6-multicast | vpnv4] }  
dont-allowas-in
```

Description

Disables EBGp from receiving and accepting a looped EBGp route from the specified neighbor, provided the number of occurrences of local AS number in AS-Path is less than or equal to the value of *as-count*.

Syntax Description

all	Specifies that the configuration change applies to all neighbors in the specified address family. If no address family is specified or if an IPv4 address is specified, the configuration change applies to all IPv4 neighbors. If an IPv6 address family is specified, the configuration change applies to all IPv6 neighbors.
<i>remoteaddr</i>	Specifies the IPv4 or IPv6 address of a <i>BGP</i> neighbor. The switch uses the IP address format to determine if the address is an IPv4 or IPv6 address.
ipv4-unicast	Specifies the IPv4 unicast address family.
ipv4-multicast	Specifies an IPv4 multicast address family.
ipv6-unicast	Specifies the IPv6 unicast address family.
ipv6-multicast	Specifies an IPv6 multicast address family.
vpn4	Specifies the VPNv4 address family for Layer 3 VPN support.

Default

This feature is disabled by default.

If no as-count is specified, the as-count defaults to 3.

If no address family is specified and an IPv4 address is detected, IPv4 unicast is the default address family.

Usage Guidelines

In a hub and spoke configuration, it becomes necessary to accept an inbound EBGp route even though the route's AS-Path contains the receiver's own AS-number. In such network topologies, this feature can be enabled.



Note

A looped AS path is always allowed for IBGP, irrespective of the BGP configuration.

All EBGp routes with looped AS-Path are silently discarded by default.

This command applies to the current VR or VRF context.

If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.



Note

For an IPv6 peer, an IPv6 address family must be specified, because an IPv6 peer does not support IPv4 address families. If no address family is specified for an IPv6 peer, the default, the IPv4 unicast address family, applies and the command fails. Similarly an IPv4 peer only supports IPv4 address families and the command fails if an IPv6 address family is specified.

To configure this feature on Layer 3 VPNs, you must configure this feature in the context of the [MPLS](#)-enabled VR; this feature is not supported for BGP neighbors on the CE (VRF) side of the PE router.

History

This command was first available in ExtremeXOS 12.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Support for Layer 3 VPNs was added in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure bgp neighbor maximum-prefix

```
configure bgp neighbor [remoteaddr | all] {address-family [ipv4-unicast
| ipv4-multicast | ipv6-unicast | ipv6-multicast | vpn4] } 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

<i>remoteaddr</i>	Specifies the IPv4 or IPv6 address of a BGP neighbor. The switch uses the IP address format to determine if the address is an IPv4 or IPv6 address.
all	Specifies that the configuration applies to all neighbors in the specified address family. If no address family is specified or if an IPv4 address is specified, the configuration applies to all IPv4 neighbors. If an IPv6 address family is specified, the configuration applies to all IPv6 neighbors.
ipv4-unicast	Specifies the IPv4 unicast address family.
ipv4-multicast	Specifies an IPv4 multicast address family.
ipv6-unicast	Specifies the IPv6 unicast address family.
ipv6-multicast	Specifies an IPv6 multicast address family.
vpn4	Specifies the VPNv4 address family for Layer 3 VPN support. This address family is applicable for PE to PE BGP neighbor sessions only. This keyword may prompt warning or error messages if executed for a regular BGP neighbor session, or for a PE to CE neighbor session.
<i>number</i>	Specifies the maximum number of prefixes that can be accepted. The range is 0 to 4294967294. A value of 0 disables prefix limit feature.

<i>percent</i>	Specifies the percentage of the maximum prefix (threshold) at which a warning message is printed in the log (and console), and/or a trap is sent to the <i>SNMP (Simple Network Management Protocol)</i> manager.
teardown	Specifies that the peer session is torn down when the maximum is exceeded.
<i>seconds</i>	Specifies the length of time before the session is re-established, if the session is torn down due to maximum prefix exceeded. If the hold-down interval is zero or not specified, it is kept down until the peer is enabled. The range is 30 to 86400 seconds.
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

This command applies to the current VR or VRF context.

If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.



Note

For an IPv6 peer, an IPv6 address family must be specified, because an IPv6 peer does not support IPv4 address families. If no address family is specified for an IPv6 peer, the default, the IPv4 unicast address family, applies and the command fails. Similarly an IPv4 peer only supports IPv4 address families and the command fails if an IPv6 address family is specified.

Configure the peer group before configuring the neighbors. To configure the peer group, use the following command:

```
configure bgp peer-group peer-group-name {address-family [ipv4-unicast | ipv4-multicast | ipv6-unicast | ipv6-multicast | vpnv4]} maximum-prefix number {{threshold percent} {teardown {holddown-interval seconds}} {send-traps}
```

To configure this feature on Layer 3 VPNs, you must configure this feature in the context of the *MPLS*-enabled VR; this feature is not supported for BGP neighbors on the CE (VRF) side of the PE router.

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:

```
configure bgp neighbor all maximum-prefix 5000 threshold 60 send-traps
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Support for Layer 3 VPNs was added in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure bgp neighbor next-hop-self

```
configure bgp neighbor [remoteaddr | all] {address-family [ipv4-unicast
| ipv4-multicast | ipv6-unicast | ipv6-multicast | vpn4] } [next-hop-
self | no-next-hop-self]
```

Description

Configures the next hop address used in the outgoing updates to be the address of the [BGP](#) connection originating the update.

Syntax Description

<i>remoteaddr</i>	Specifies an IP address.
all	Specifies all neighbors.
ipv4-unicast	Specifies the IPv4 unicast address family.
ipv4-multicast	Specifies an IPv4 multicast address family.
ipv6-unicast	Specifies the IPv6 unicast address family.
ipv6-multicast	Specifies an IPv6 multicast address family.
vpn4	Specifies the VPNv4 address family for Layer 3 VPN support. This address family is applicable for PE to PE BGP neighbor sessions only. This keyword may prompt warning or error messages if executed for a regular BGP neighbor session, or for a PE to CE neighbor session.

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 (lets BGP decide what would be the next hop).

Default

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

This command applies to the current VR or VRF context. These settings apply to the peer group and all neighbors of the peer group.



Note

The BGP neighbor must be disabled before you can change the configuration with this command.

If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.



Note

For an IPv6 peer, an IPv6 address family must be specified, because an IPv6 peer does not support IPv4 address families. If no address family is specified for an IPv6 peer, the default, the IPv4 unicast address family, applies and the command fails. Similarly an IPv4 peer only supports IPv4 address families and the command fails if an IPv6 address family is specified.

To configure this feature on Layer 3 VPNs, you must configure this feature in the context of the MPLS-enabled VR; this feature is not supported for BGP neighbors on the CE (VRF) side of the PE router.

Example

The following command configures the next hop address used in the updates to be the address of the BGP connection originating it:

```
configure bgp neighbor 172.16.5.25 next-hop-self
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for Layer 3 VPNs was added in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure bgp neighbor no-dampening

```
configure bgp neighbor [remoteaddr | all] {address-family [ipv4-unicast
| ipv4-multicast | ipv6-unicast | ipv6-multicast | vpnv4] } no-
dampening
```

Description

Configures no route flap dampening over *BGP* peer sessions (disables route flap dampening).

Syntax Description

<i>remoteaddr</i>	Specifies the IPv4 or IPv6 address of a BGP neighbor. The switch uses the IP address format to determine if the address is an IPv4 or IPv6 address.
all	Specifies that the configuration applies to all neighbors in the specified address family. If no address family is specified or if an IPv4 address is specified, the configuration applies to all IPv4 neighbors. If an IPv6 address family is specified, the configuration applies to all IPv6 neighbors.
ipv4-unicast	Specifies the IPv4 unicast address family.
ipv4-multicast	Specifies an IPv4 multicast address family.
ipv6-unicast	Specifies the IPv6 unicast address family.
ipv6-multicast	Specifies an IPv6 multicast address family.
vpnv4	Specifies the VPNv4 address family for Layer 3 VPN support.

Default

This feature is disabled by default.

If no address family is specified and an IPv4 address is detected, IPv4 unicast is the default address family.

Usage Guidelines

If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.



Note

For an IPv6 peer, an IPv6 address family must be specified, because an IPv6 peer does not support IPv4 address families. If no address family is specified for an IPv6 peer, the default, the IPv4 unicast address family, applies and the command fails. Similarly an IPv4 peer only supports IPv4 address families and the command fails if an IPv6 address family is specified.

Use the following command to enable route flap dampening for BGP neighbors:

```
configure bgp neighbor [all | remoteaddr] {address-family [ipv4-unicast
| ipv4-multicast | ipv6-unicast | ipv6-multicast | vpn4]} dampening
{{half-life half-life-minutes {reuse-limit reuse-limit-number suppress-
limit suppress-limit-number max-suppress max-suppress-minutes} | policy-
filter [policy-name | none]}
```

To configure this feature on Layer 3 VPNs, you must configure this feature in the context of the *MPLS*-enabled VR; this feature is not supported for BGP neighbors on the CE (VRF) side of the PE router.

Example

The following command disables route flap dampening to the BGP neighbor at 192.168.1.22:

```
configure bgp neighbor 192.168.1.22 no-dampening
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Support for Layer 3 VPNs was added in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure bgp neighbor password

```
configure bgp neighbor [all | remoteaddr] password [none | {encrypted}
tcpPassword]
```

Description

Configures an RSA Data Security, Inc. *MD5 (Message-Digest algorithm 5)* Message-Digest Algorithm secret password for a neighbor.

Syntax Description

all	Specifies all IPv4 and IPv6 neighbors.
<i>remoteaddr</i>	Specifies the IPv4 or IPv6 address of a <i>BGP</i> neighbor.
none	Specifies not to use a password
encrypted	Specifies an encrypted string; do not use.
<i>tcpPassword</i>	Specifies a password string.

Default

N/A.

Usage Guidelines

This command applies to the current VR or VRF context.

You must disable the BGP neighbor before changing the password.

When a password is configured, TCP RSA Data Security, Inc. MD5 Message-Digest Algorithm 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.

To change any one of the following parameters you must disable and re-enable the peer session:

- timer
- source-interface
- soft-in-reset
- password

Changing a route reflector client automatically disables and enables the peer session.

The encrypted option is used by the switch when generating a configuration file, and when parsing a switch-generated configuration file. Do not select the encrypted option in the CLI.

Example

The following command configures the password for a neighbor as Extreme:

```
configure bgp neighbor 192.168.1.5 password extreme
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure bgp neighbor peer-group

```
configure bgp neighbor [all | remoteaddr] peer-group [peer-group-name |  
none] {acquire-all}
```

Description

Configures an existing neighbor as the member of a peer group.

Syntax Description

all	Specifies all neighbors.
<i>remoteaddr</i>	Specifies the IPv4 or IPv6 address of a BGP neighbor.
<i>peer-group-name</i>	Specifies a peer group name.
none	Removes the neighbor from the peer group.
acquire-all	Specifies that all parameters should be inherited by the neighbor from the peer group.

Default

By default, remote AS (if configured for the peer group), source-interface, outbound route policy, send-community and next-hop-self settings are inherited.

Usage Guidelines

This command applies to the current VR or VRF context.

If acquire-all is not specified, only the default parameters are inherited by the neighbor.

When you remove a neighbor from a peer group, it retains the parameter settings of the group. The parameter values are not reset to those the neighbor had before it inherited the peer group values.

To create a new neighbor and add it to a BGP peer group, use the following command:

```
create bgp neighbor remoteaddr peer-group peer-group-name {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 a remote AS configured.

If you are adding an IPv4 peer to a peer group and no IPv4 address family capabilities are assigned to the specified peer group, the IPv4 unicast and multicast address families are automatically enabled for that peer group. If you adding an IPv6 peer to a peer group and no IPv6 address family capabilities are assigned to the peer group, you must explicitly enable the IPv6 address family capabilities you want to support.



Note

If the peer group or any member of the peer group has been configured with an IPv4 or IPv6 address family, the peer group only accepts peers that are configured to use that family. For example, if a peer group is configured for the IPv4 unicast address family, the switch will not allow you to add an IPv6 peer. Likewise, an IPv6 peer group cannot accept an IPv4 peer.

Example

The following command configures an existing neighbor as the member of the peer group outer:

```
configure bgp neighbor 192.1.1.22 peer-group outer
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure bgp neighbor route-policy

```
configure bgp neighbor [remoteaddr | all] {address-family [ipv4-unicast  
| ipv4-multicast | ipv6-unicast | ipv6-multicast | vpnv4]} route-policy  
[in | out] [none | policy]
```

Description

Configures a route map filter for a neighbor.

Syntax Description

<i>remoteaddr</i>	Specifies the IPv4 or IPv6 address of a <i>BGP</i> neighbor. The switch uses the IP address format to determine if the address is an IPv4 or IPv6 address.
all	Specifies that the configuration applies to all neighbors in the specified address family. If no address family is specified or if an IPv4 address is specified, the configuration applies to all IPv4 neighbors. If an IPv6 address family is specified, the configuration applies to all IPv6 neighbors.
ipv4-unicast	Specifies the IPv4 unicast address family.
ipv4-multicast	Specifies an IPv4 multicast address family.
ipv6-unicast	Specifies the IPv6 unicast address family.
ipv6-multicast	Specifies an IPv6 multicast address family.
vpnv4	Specifies the VPNv4 address family for Layer 3 VPN support. This address family is applicable for PE to PE BGP neighbor sessions only. This keyword may prompt warning or error messages if executed for a regular BGP neighbor session, or for a PE to CE neighbor session.
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.
<i>policy</i>	Specifies a policy.

Default

If no address family is specified and an IPv4 address is detected, IPv4 unicast is the default address family.

Usage Guidelines

This command applies to the current VR or VRF context.

The policy can be installed on the input or output side of the router. The policy is used to modify or filter the NLRI information and the path attributes associated with it when exchanging updates with the neighbor.



Note

A policy file applied to BGP neighbors cannot have NLRI for both IPv4 and IPv6 address families defined in the same policy file.

If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.

**Note**

For an IPv6 peer, an IPv6 address family must be specified, because an IPv6 peer does not support IPv4 address families. If no address family is specified for an IPv6 peer, the default, the IPv4 unicast address family, applies and the command fails. Similarly an IPv4 peer only supports IPv4 address families and the command fails if an IPv6 address family is specified.

To configure this feature on Layer 3 VPNs, you must configure this feature in the context of the *MPLS*-enabled VR; this feature is not supported for BGP neighbors on the CE (VRF) side of the PE router.

Example

The following command configures the route-policy filter for a neighbor based on the policy nosales:

```
configure bgp neighbor 192.168.1.22 route-policy in nosales
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Support for Layer 3 VPNs was added in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure bgp neighbor route-reflector-client

```
configure bgp neighbor [remoteaddr | all] [route-reflector-client | no-  
route-reflector-client]
```

Description

Configures a *BGP* neighbor to be a route reflector client.

Syntax Description

<i>remoteaddr</i>	Specifies the IPv4 or IPv6 address of a BGP neighbor.
all	Specifies all IPv4 and IPv6 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.

This command applies to the current VR or VRF context.

When changing the route reflector status of a peer, the peer is automatically disabled and re-enabled and a warning message appears on the console and in the log.

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:

```
configure bgp neighbor 192.168.1.5 route-reflector-client
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure bgp neighbor send-community

```
configure bgp neighbor [remoteaddr | all] {address-family [ipv4-unicast
| ipv4-multicast | ipv6-unicast | ipv6-multicast | vpnv4]} [send-
community | dont-send-community] {both | extended | standard}
```

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

<i>remoteaddr</i>	Specifies an IP address of a BGP neighbor.
all	Specifies all neighbors.
ipv4-unicast	Specifies the IPv4 unicast address family.
ipv4-multicast	Specifies an IPv4 multicast address family.
ipv6-unicast	Specifies the IPv6 unicast address family.
ipv6-multicast	Specifies an IPv6 multicast address family.
vpnv4	Specifies the VPNv4 address family for Layer 3 VPN support. This address family is applicable for PE to PE BGP neighbor sessions only. This keyword may prompt warning or error messages if executed for a regular BGP neighbor session, or for a PE to CE neighbor session.
send-community	Specifies to include the community path attribute.
dont-send-community	Specifies not to include the community path attribute.
both	Send both standard and extended community attributes to this BGP neighbor, or neighbors in peer group
extended	Send only extended communities to this BGP neighbor or neighbors in peer group
standard	Send only standard communities to this BGP neighbor or neighbors in peer group

Default

If no address family is specified, IPv4 unicast is the default. If no optional keyword (both, standard or extended) is specified, standard is assumed.

Usage Guidelines

A BGP community is a group of BGP destinations that require common handling. ExtremeXOS supports the following well-known BGP community attributes:

- no-export
- no-advertise
- no-export-subconfed

The command is additive; that is, if the command is executed twice with the standard or extended option, both the extended and standard communities are sent to the BGP neighbor.

This command applies to the current VR or VRF context.

If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.



Note

For an IPv6 peer, an IPv6 address family must be specified, because an IPv6 peer does not support IPv4 address families. If no address family is specified for an IPv6 peer, the default, the IPv4 unicast address family, applies and the command fails. Similarly an IPv4 peer only supports IPv4 address families and the command fails if an IPv6 address family is specified.

To configure this feature on Layer 3 VPNs, you must configure this feature in the context of the *MPLS*-enabled VR; this feature is not supported for BGP neighbors on the CE (VRF) side of the PE router.

Example

The following command includes the community path attribute associated with a BGP NLRI in the route updates sent to all BGP neighbors:

```
configure bgp neighbor all send-community
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Options to control the advertisement of extended community attributes were added in ExtremeXOS 12.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Support for Layer 3 VPNs was added in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure bgp neighbor shutdown-priority

```
configure bgp neighbor [all | remoteaddr] shutdown-priority number
```

Description

Configures the shutdown priority for a *BGP* neighbor.

Syntax Description

<i>remoteaddr</i>	Specifies the IPv4 or IPv6 address of a BGP neighbor.
all	Specifies all IPv4 and IPv6 neighbors.
<i>number</i>	Specifies the shutdown priority. The range is 0 - 65,535.

Default

The default value is 1024.

Usage Guidelines



Note

This command is not currently supported, and is not recommended for use.

Higher priority values lower the chance of a BGP neighbor to be automatically disabled in case BGP or the system goes to a low memory condition.

Example

The following command configures the shutdown priority of the BGP neighbor 10.0.20.1 to 500:

```
configure bgp neighbor 10.0.20.1 shutdown-priority 1000
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure bgp neighbor soft-reset

```
configure bgp neighbor [remoteaddr | all] {address-family [ipv4-unicast
| ipv4-multicast | ipv6-unicast | ipv6-multicast | vpn4 | l2vpn-evpn] }
soft-reset {in | out}
```

Description

Applies the current input or output routing policy to the routing information already exchanged with the neighbor.

Syntax Description

<i>remoteaddr</i>	Specifies an IP address of a <i>BGP</i> neighbor.
all	Specifies all neighbors.
ipv4-unicast	Specifies the IPv4 unicast address family.
ipv4-multicast	Specifies an IPv4 multicast address family.
ipv6-unicast	Specifies the IPv6 unicast address family.
ipv6-multicast	Specifies an IPv6 multicast address family.
vpnv4	Specifies the VPNv4 address family for Layer 3 VPN support. This address family is applicable for PE to PE BGP neighbor sessions only. This keyword may prompt warning or error messages if executed for a regular BGP neighbor session, or for a PE to CE neighbor session.
l2vpn-evpn	Specifies the Layer 2 VPN-EVPN address family.
soft-reset	Do a soft reconfiguration for the BGP neighbor.
in	Specifies to apply the input routing policy.
out	Specifies to apply the output routing policy.

Default

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

The input/output policy is determined by the route policy configured for the neighbor on the input and/or output side of the router. This command does not affect the switch configuration.

If both the local BGP neighbor and the neighbor router support the route refresh capability, a dynamic soft input reset can be performed. The `configure bgp neighbor soft-reset` command triggers the generation of a Route-Refresh message to the neighbor. As a response to the Route-Refresh message, the neighbor sends the entire BGP routing table in updates and the switch applies the appropriate routing policy to the updates.

This command applies to the current VR or VRF context.

If the route-refresh capability is not supported by the neighbor, the `configure bgp neighbor soft-reset` command reprocesses the BGP route database using the policy configured for that neighbor.

If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.

**Note**

For an IPv6 peer, an IPv6 address family must be specified, because an IPv6 peer does not support IPv4 address families. If no address family is specified for an IPv6 peer, the default, the IPv4 unicast address family, applies and the command fails. Similarly an IPv4 peer only supports IPv4 address families and the command fails if an IPv6 address family is specified.

To configure this feature on Layer 3 VPNs, you must configure this feature in the context of the *MPLS*-enabled VR; this feature is not supported for BGP neighbors on the CE (VRF) side of the PE router.

Example

The following command applies the current input routing policy to the routing information already exchanged with the neighbor:

```
# configure bgp neighbor 192.168.1.5 soft-reset in
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for Layer 3 VPNs was added in ExtremeXOS 15.3.

Support for Layer 2 VPN-EVPN was added in ExtremeXOS 30.5.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure bgp neighbor source-interface

```
configure bgp neighbor [remoteaddr | all] source-interface [any |  
  ipaddress ipAddr]
```

Description

Changes the *BGP* source interface for TCP connections.

Syntax Description

<i>remoteaddr</i>	Specifies the IPv4 or IPv6 address of a BGP neighbor.
all	Specifies all IPv4 and IPv6 neighbors.
any	Specifies any source interface.
<i>ipAddr</i>	Specifies the IP address of a source interface.

Default

Any.

Usage Guidelines

The source interface IP address must be a valid IP address of any *VLAN* configured on the switch.

This command applies to the current VR or VRF context.

Example

The following command changes the BGP source interface to 10.43.55.10:

```
configure bgp neighbor 192.168.1.5 source-interface ipaddress 10.43.55.10
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure bgp neighbor timer

```
configure bgp neighbor [remoteaddr | all] timer keep-alive keepalive  
hold-time holdtime
```

Description

Configures the *BGP* neighbor timers.

Syntax Description

<i>remoteaddr</i>	Specifies the IPv4 or IPv6 address of a BGP neighbor.
all	Specifies all IPv4 and IPv6 neighbors.
<i>keepalive</i>	Specifies a BGP neighbor timer keepalive time in seconds. The range is 0 to 21,845 seconds.
<i>holdtime</i>	Specifies a BGP neighbor timer hold time in seconds. The range is 0 and 3to65,535 seconds.

Default

The default keepalive setting is 60 seconds. The default hold time is 180 seconds.

Usage Guidelines

You must disable the BGP neighbor before changing the timer values.

This command applies to the current VR or VRF context.

Example

The following command configures the BGP neighbor timers:

```
configure bgp neighbor 192.168.1.5 timer keep-alive 120 hold-time 360
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure bgp neighbor weight

```
configure bgp neighbor [remoteaddr | all] weight weight
```

Description

Assigns a locally-used weight to a neighbor connection for the route selection algorithm.

Syntax Description

<i>remoteaddr</i>	Specifies the IPv4 or IPv6 address of a BGP neighbor.
all	Specifies all IPv4 and IPv6 neighbors.
<i>weight</i>	Specifies a BGP neighbor weight.

Default

By default, the weight is 1.

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 65,535.

BGP selects routes based on the following precedence (from highest to lowest):

- higher weight
- higher local preference
- shortest length (shortest AS path)
- lowest origin code
- lowest MED
- route from external peer
- lowest cost to Next Hop
- lowest routerID

This command applies to the current VR or VRF context.

Example

The following command assigns a locally used weight of 10 to a neighbor connection:

```
configure bgp neighbor 192.168.1.5 weight 10
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure bgp peer-group allowas-in

```
configure bgp peer-group peer-group-name {address-family [ipv4-unicast
| ipv4-multicast | ipv6-unicast | ipv6-multicast | vpnv4] } allowas-in
{max-as-occurrence as-count}
```

Description

Configures *BGP* to receive and accept a looped BGP route from the neighbors of the specified peer group, provided the number of occurrences of local AS number in AS-Path is less than or equal to that specified in *as-count*.

Syntax Description

<i>peer-group-name</i>	Specifies a peer group.
ipv4-unicast	Specifies the IPv4 unicast address family.
ipv4-multicast	Specifies an IPv4 multicast address family.
ipv6-unicast	Specifies the IPv6 unicast address family.
ipv6-multicast	Specifies an IPv6 multicast address family.
vpnv4	Specifies the VPNv4 address family for Layer 3 VPN support.
<i>as-count</i>	The maximum number of occurrences of local AS number in the received route AS-Path. If the number of occurrences of local AS number in AS-Path is more than <i>as-count</i> , the route is not accepted. The valid range is from 1-16.

Default

This feature is disabled by default.

If no *as-count* is specified, the *as-count* defaults to 3.

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

In a hub and spoke configuration, it becomes necessary to accept an inbound BGP route even though the route's AS-Path contains the receiver's own AS-number. In such network topologies, this feature can be enabled.

This feature can also be enabled for both IBGP and EBGP neighbors, wherever necessary.

This command applies to the current VR or VRF context.



Note

BGP neighbors do not inherit the *allowas-in* configuration from their peer group unless you explicitly specify the *acquire-all* option when adding a neighbor to a peer-group.

If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.

**Note**

If the specified peer group contains IPv6 peers, it is an IPv6 peer group and you must specify an IPv6 address-family. When the specified peer group is an IPv6 peer group, this command fails if no address family is specified or if an IPv4 address-family is specified. This command also fails if an IPv6 address family is specified for an IPv4 peer-group.

To configure this feature on Layer 3 VPNs, you must configure this feature in the context of the MPLS-enabled VR; this feature is not supported for BGP neighbors on the CE (VRF) side of the PE router.

Example

The following example enables BGP to accept looped BGP routes that contains a maximum of 8 occurrences of receiver's AS-number in AS-Path attribute:

```
configure bgp peer-group internal allowas-in max-as-occurrence 8
```

History

This command was first available in ExtremeXOS 12.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Support for Layer 3 VPNs was added in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure bgp peer-group connect-retry

```
configure bgp peer-group peer-group-name connect-retry seconds
```

Description

Configures the BGP retry timer of the specified peer group.

Syntax Description

<i>peer-group-name</i>	Specifies a peer group.
connect-retry	Specifies to configure the time interval between attempts to establish a TCP connection to a configured peer.
<i>seconds</i>	Specifies the retry time in seconds. Default is 30 seconds.

Default

This feature is disabled by default.

Usage Guidelines

Example

The following command configures the BGP peer group "internal" retry timer to 120 seconds:

```
configure bgp peer-group internal connect-retry 120
```

History

This command was first available in ExtremeXOS 31.6.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure bgp peer-group dampening

```
configure bgp peer-group peer-group-name {address-family [ipv4-unicast | ipv4-multicast | ipv6-unicast | ipv6-multicast | vpn4]}
dampening {{half-life half-life-minutes {reuse-limit reuse-limit-number supress-limit suppress-limit-number max-suppress max-suppress-minutes}} | policy-filter [policy-name | none]}
```

Description

Configures route flap dampening for a BGP peer group.

Syntax Description

<i>peer-group-name</i>	Specifies a peer group.
ipv4-unicast	Specifies the IPv4 unicast address family.
ipv4-multicast	Specifies an IPv4 multicast address family.

ipv6-unicast	Specifies the IPv6 unicast address family.
ipv6-multicast	Specifies an IPv6 multicast address family.
vpnv4	Specifies the VPNv4 address family for Layer 3 VPN support.
<i>half-life-minutes</i>	Specifies the dampening half life.
<i>reuse-limit-number</i>	Specifies the reuse limit.
<i>suppress-limit-number</i>	Specifies the suppress limit.
<i>max-suppress-minutes</i>	Specifies the maximum hold down time.
<i>policy-name</i>	Specifies a policy.
none	Removes any policy association.

Default

This feature is disabled by default.

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

This command applies to the current VR or VRF context.

The half life is the period of time, in minutes, during which the accumulated penalty of a route is reduced by half. The range is 1 to 45 minutes, and the default is 15 minutes.

The reuse limit is the penalty value below which a route is used again. The range is 1-20,000, and the default is 750.

The suppress limit is the penalty value above which a route is suppressed. The range is 1-20,000, and the default is 2,000.

The maximum hold down time is the maximum time a route can be suppressed, no matter how unstable it has been, as long as it no longer flaps. The range is 1-255 minutes, and the default is 4 * the half life.

If you change dampening parameters when routes are in suppressed or history state, the new dampening parameters apply only to routes in the active state. Routes in the suppressed or history state continue to use the old dampening parameters until they become active, at which time they use the updated dampening parameters.

Instead of explicitly configuring the dampening parameters using the command line, you can associate a policy using the policy-filter option. Multiple sets of parameters can be supplied using a policy.

Use the following command to disable route flap dampening for a BGP peer-group:

```
configure bgp peer-group peer-group-name {address-family [ipv4-unicast | ipv4-multicast | ipv6-unicast | ipv6-multicast | vpnv4]} no-dampening
```

If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.

**Note**

If the specified peer group contains IPv6 peers, it is an IPv6 peer group and you must specify an IPv6 address-family. When the specified peer group is an IPv6 peer group, this command fails if no address family is specified or if an IPv4 address-family is specified. This command also fails if an IPv6 address family is specified for an IPv4 peer-group.

To configure this feature on Layer 3 VPNs, you must configure this feature in the context of the MPLS-enabled VR; this feature is not supported for BGP neighbors on the CE (VRF) side of the PE router.

Example

The following command configures route flap dampening for the BGP peer group outer:

```
configure bgp peer-group outer dampening
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Support for Layer 3 VPNs was added in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure bgp peer-group dont-allowas-in

```
configure bgp peer-group peer-group-name {address-family [ipv4-unicast |  
ipv4-multicast | ipv6-unicast | ipv6-multicast | vpn4] } dont-allowas-  
in
```

Description

Disables BGP from receiving and accepting a looped BGP route from the neighbors of the specified peer group, provided the number of occurrences of local AS number in AS-Path is less than or equal to that specified in `as-count`.

Syntax Description

<i>peer-group-name</i>	Specifies a peer group.
ipv4-unicast	Specifies the IPv4 unicast address family.
ipv4-multicast	Specifies an IPv4 multicast address family.
ipv6-unicast	Specifies the IPv6 unicast address family.
ipv6-multicast	Specifies an IPv6 multicast address family.
vpnv4	Specifies the VPNv4 address family for Layer 3 VPN support.

Default

This feature is disabled by default.

If no as-count is specified, the as-count defaults to 3.

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

This command applies to the current VR or VRF context.

If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.



Note

If the specified peer group contains IPv6 peers, it is an IPv6 peer group and you must specify an IPv6 address-family. When the specified peer group is an IPv6 peer group, this command fails if no address family is specified or if an IPv4 address-family is specified. This command also fails if an IPv6 address family is specified for an IPv4 peer-group.

To configure this feature on Layer 3 VPNs, you must configure this feature in the context of the MPLS-enabled VR; this feature is not supported for BGP neighbors on the CE (VRF) side of the PE router.



Note

BGP neighbors do not inherit the allow-as-in configuration from their peer group unless you explicitly specify the acquire-all option when adding a neighbor to a peer-group.

History

This command was first available in ExtremeXOS 12.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Support for Layer 3 VPNs was added in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure bgp peer-group maximum-prefix

```
configure bgp peer-group peer-group-name {address-family [ipv4-unicast |
ipv4-multicast | ipv6-unicast | ipv6-multicast | vpnv4]} maximum-prefix
number {{threshold percent} {teardown {holddown-interval seconds}}
{send-traps}
```

Description

Configures the maximum number of IP prefixes accepted for all neighbors in the peer group.

Syntax Description

name	Specifies a peer group.
ipv4-unicast	Specifies the IPv4 unicast address family.
ipv4-multicast	Specifies an IPv4 multicast address family.
ipv6-unicast	Specifies the IPv6 unicast address family.
ipv6-multicast	Specifies an IPv6 multicast address family.
vpnv4	Specifies the VPNv4 address family for Layer 3 VPN support.
<i>number</i>	Specifies the maximum number of prefixes that can be accepted. The range is 0 to 4294967294. A value of 0 disables prefix limit feature.
<i>percent</i>	Specifies the percentage of the maximum prefix (threshold) at which a warning message is printed in the log (and on the console). An <i>SNMP</i> trap can also be sent.
teardown	Specifies that the peer session is torn down when the maximum is exceeded.
<i>seconds</i>	Specifies the length of time before the session is re-established, if the session has been torn down due to exceeding the max limit. If the hold down interval is 0 or not specified, it is kept down until the peer is enabled. The range is 30 to 86400 seconds.
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.

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

This command applies to the current VR or VRF context.

If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.



Note

If the specified peer group contains IPv6 peers, it is an IPv6 peer group and you must specify an IPv6 address-family. When the specified peer group is an IPv6 peer group, this command fails if no address family is specified or if an IPv4 address-family is specified. This command also fails if an IPv6 address family is specified for an IPv4 peer-group.

Configure the peer group before configuring the neighbors. To configure the neighbors, use the following command:

```
configure bgp neighbor 192.168.1.1 maximum-prefix
```

After you enter this command, the switch automatically disables and enables all neighbors in the peer group before the change takes effect.

To configure this feature on Layer 3 VPNs, you must configure this feature in the context of the MPLS-enabled VR; this feature is not supported for BGP neighbors on the CE (VRF) side of the PE router.

Example

The following command configures the maximum number of IP prefixes accepted from the peer group outer to 5000, sets the threshold for warning messages to 60%, and specifies SNMP traps:

```
configure bgp peer-group outer maximum-prefix 5000 threshold 60 send-traps
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Support for Layer 3 VPNs was added in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure bgp peer-group next-hop-self

```
configure bgp peer-group peer-group-name {address-family [ipv4-unicast |
ipv4-multicast | ipv6-unicast | ipv6-multicast | vpnv4] } [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

<i>peer-group-name</i>	Specifies a peer group.
ipv4-unicast	Specifies the IPv4 unicast address family.
ipv4-multicast	Specifies an IPv4 multicast address family.
ipv6-unicast	Specifies the IPv6 unicast address family.
ipv6-multicast	Specifies an IPv6 multicast address family.
vpnv4	Specifies the VPNv4 address family for Layer 3 VPN support.
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 (Let the BGP protocol decide the next hop).

Default

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

This command applies to the current VR or VRF context.

These settings apply to the peer group and all neighbors of the peer group.

If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.

**Note**

If the specified peer group contains IPv6 peers, it is an IPv6 peer group and you must specify an IPv6 address-family. When the specified peer group is an IPv6 peer group, this command fails if no address family is specified or if an IPv4 address-family is specified. This command also fails if an IPv6 address family is specified for an IPv4 peer-group.

To configure this feature on Layer 3 VPNs, you must configure this feature in the context of the *MPLS*-enabled VR; this feature is not supported for BGP neighbors on the CE (VRF) side of the PE router.

Example

The following command configures the next hop address used in the updates to be the address of the BGP connection originating it:

```
configure bgp peer-group outer next-hop-self
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Support for Layer 3 VPNs was added in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure bgp peer-group no-dampening

```
configure bgp peer-group peer-group-name {address-family [ipv4-unicast |  
ipv4-multicast | ipv6-unicast | ipv6-multicast | vpn4] } no-dampening
```

Description

Configures no route flap dampening for a *BGP* peer group (disables route flap dampening).

Syntax Description

<i>peer-group-name</i>	Specifies a BGP peer group.
ipv4-unicast	Specifies the IPv4 unicast address family.
ipv4-multicast	Specifies an IPv4 multicast address family.
ipv6-unicast	Specifies the IPv6 unicast address family.
ipv6-multicast	Specifies an IPv6 multicast address family.
vpnv4	Specifies the VPNv4 address family for Layer 3 VPN support.

Default

This feature is disabled by default.

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.



Note

If the specified peer group contains IPv6 peers, it is an IPv6 peer group and you must specify an IPv6 address-family. When the specified peer group is an IPv6 peer group, this command fails if no address family is specified or if an IPv4 address-family is specified. This command also fails if an IPv6 address family is specified for an IPv4 peer-group.

Use the following command to enable route flap dampening for a BGP peer-group:

```
configure bgp peer-group peer-group-name {address-family [ipv4-unicast
| ipv4-multicast | ipv6-unicast | ipv6-multicast | vpnv4]} dampening
{{half-life half-life-minutes {reuse-limit reuse-limit-number suppress-
limit suppress-limit-number max-suppress max-suppress-minutes}} |
policy-filter [policy-name | none]}
```

To configure this feature on Layer 3 VPNs, you must configure this feature in the context of the *MPLS*-enabled VR; this feature is not supported for BGP neighbors on the CE (VRF) side of the PE router.

Example

The following command disables route flap dampening to the BGP peer group outer:

```
configure bgp peer-group outer no-dampening
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Support for Layer 3 VPNs was added in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure bgp peer-group password

```
configure bgp peer-group peer-group-name password [none | {encrypted}  
            tcpPassword]
```

Description

Configures the TCP RSA Data Security, Inc. [MD5](#) Message-Digest Algorithm secret password for a peer group and all neighbors of the peer group.

Syntax Description

<i>peer-group-name</i>	Specifies a peer group.
none	Specifies no password.
<i>tcpPassword</i>	Specifies a password.
encrypted	Specifies an encrypted string.

Default

N/A.

Usage Guidelines

Changes made to the parameters of a peer group are applied to all neighbors in the peer group.

This command applies to the current VR or VRF context.

After you enter this command, the switch automatically disables and enables all neighbors in the peer group before the change takes effect.

Example

The following command configures the password as Extreme for the peer group outer and its neighbors:

```
configure bgp peer-group outer password extreme
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the *BGP* feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure bgp peer-group remote-AS-number

```
configure bgp peer-group peer-group-name remote-AS-number number
```

Description

Configures the remote AS number for a peer group and all the neighbors of the peer group.

Syntax Description

<i>peer-group-name</i>	Specifies a peer group.
<i>number</i>	Specifies a remote AS number. The range is 1 to 4294967295.

Default

N/A.

Usage Guidelines

The AS number is a 4-byte AS number in either the ASPLAIN or the ASDOT format as described in RFC 5396, Textual Representation of Autonomous System (AS) Numbers.

Changes made to the parameters of a peer group are applied to all neighbors in the peer group.

After you enter this command, the switch automatically disables and enables all neighbors in the peer group before the change takes effect.

Example

The following example configures the remote AS number for the peer group outer and its neighbors using the ASPLAIN 4-byte AS number format:

```
configure bgp peer-group outer remote-AS-number 65536
```

The following example configures the remote AS number for the peer group abc and its neighbors using the ASDOT 4-byte AS number format:

```
configure bgp peer-group abc remote-AS-number 1.10
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for 4-byte AS numbers was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the *BGP* feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure bgp peer-group route-policy

```
configure bgp peer-group peer-group-name {address-family [ipv4-unicast  
| ipv4-multicast | ipv6-unicast | ipv6-multicast | vpn4] } route-policy  
[in | out] [none | policy]
```

Description

Configures the policy for a peer group and all the neighbors of the peer group.

Syntax Description

<i>peer-group-name</i>	Specifies a peer group.
ipv4-unicast	Specifies the IPv4 unicast address family.
ipv4-multicast	Specifies an IPv4 multicast address family.
ipv6-unicast	Specifies the IPv6 unicast address family.
ipv6-multicast	Specifies an IPv6 multicast address family.
vpn4	Specifies the VPNv4 address family for Layer 3 VPN support.
in	Specifies to install the policy on the input side.
out	Specifies to install the policy on the output side.
none	Specifies to remove the filter.
<i>policy</i>	Specifies a policy.

Default

There is no default policy configuration.

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

This command applies to the current VR or VRF context.

If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.

**Note**

If the specified peer group contains IPv6 peers, it is an IPv6 peer group and you must specify an IPv6 address-family. When the specified peer group is an IPv6 peer group, this command fails if no address family is specified or if an IPv4 address-family is specified. This command also fails if an IPv6 address family is specified for an IPv4 peer-group.

To configure this feature on Layer 3 VPNs, you must configure this feature in the context of the *MPLS*-enabled VR; this feature is not supported for *BGP* neighbors on the CE (VRF) side of the PE router.

Example

The following command configures the route policy for the peer group outer and its neighbors using the policy nosales:

```
configure bgp peer-group outer route-policy in nosales
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Support for Layer 3 VPNs was added in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure bgp peer-group route-reflector-client

```
configure bgp peer-group peer-group-name [route-reflector-client | no-  
route-reflector-client]
```

Description

Configures all the peers in a peer group to be a route reflector client.

Syntax Description

<i>peer-group-name</i>	Specifies a peer group.
route-reflector-client	Specifies that all the neighbors in the peer group be a route reflector client.
no-route-reflector-client	Specifies that all the neighbors in 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.

This command applies to the current VR or VRF context.

The peer group must be in the same AS of this router.

Example

The following command configures the peer group outer as a route reflector client:

```
configure bgp peer-group outer route-reflector-client
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the *BGP* feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure bgp peer-group send-community

```
configure bgp peer-group peer-group-name {address-family [ipv4-unicast  
| ipv4-multicast | ipv6-unicast | ipv6-multicast | vpn4]} [send-  
community | dont-send-community] {both | extended | standard}
```

Description

Configures whether communities should be sent to neighbors as part of route updates.

Syntax Description

<i>peer-group-name</i>	Specifies a peer group.
ipv4-unicast	Specifies the IPv4 unicast address family.
ipv4-multicast	Specifies an IPv4 multicast address family.
ipv6-unicast	Specifies the IPv6 unicast address family.
ipv6-multicast	Specifies an IPv6 multicast address family.
vpnv4	Specifies the VPNv4 address family for Layer 3 VPN support.
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.
both	Send both standard and extended community attributes to this <i>BGP</i> neighbor, or neighbors in peer group.
extended	Send only extended communities to this BGP neighbor or neighbors in peer group.
standard	Send only standard communities to this BGP neighbor or neighbors in peer group.

Default

If no address family is specified, IPv4 unicast is the default. If no optional keyword (both, standard or extended) is specified, standard is assumed.

Usage Guidelines

These settings apply to the peer group and all neighbors of the peer group.

This command applies to the current VR or VRF context.

If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.



Note

If the specified peer group contains IPv6 peers, it is an IPv6 peer group and you must specify an IPv6 address-family. When the specified peer group is an IPv6 peer group, this command fails if no address family is specified or if an IPv4 address-family is specified. This command also fails if an IPv6 address family is specified for an IPv4 peer-group.

The command is additive; that is, if the command is executed twice with the standard or extended option, both the extended and standard communities are sent to the BGP neighbor.

To configure this feature on Layer 3 VPNs, you must configure this feature in the context of the *MPLS*-enabled VR; this feature is not supported for BGP neighbors on the CE (VRF) side of the PE router.

Example

The following command configures communities to be sent to neighbors as part of route updates:

```
configure bgp peer-group outer send-community
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Options to control the advertisement of extended community attributes were added in ExtremeXOS 12.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Support for Layer 3 VPNs was added in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure bgp peer-group soft-reset

```
configure bgp peer-group peer-group-name {address-family [ipv4-unicast |  
ipv4-multicast | ipv6-unicast | ipv6-multicast | vpn4]} soft-reset {in  
| out}
```

Description

Applies the current input/output routing policy to the neighbors in the peer group.

Syntax Description

<i>peer-group-name</i>	Specifies a peer group.
ipv4-unicast	Specifies the IPv4 unicast address family.
ipv4-multicast	Specifies an IPv4 multicast address family.
ipv6-unicast	Specifies the IPv6 unicast address family.
ipv6-multicast	Specifies an IPv6 multicast address family.
vpn4	Specifies the VPNv4 address family for Layer 3 VPN support.
in	Specifies to apply the input routing policy.
out	Specifies to apply the output routing policy.

Default

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

The input/output routing policy is determined by the route policy 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.

This command applies to the current VR or VRF context.

If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.



Note

If the specified peer group contains IPv6 peers, it is an IPv6 peer group and you must specify an IPv6 address-family. When the specified peer group is an IPv6 peer group, this command fails if no address family is specified or if an IPv4 address-family is specified. This command also fails if an IPv6 address family is specified for an IPv4 peer-group.

Changes made to the parameters of a peer group are applied to all neighbors in the peer group.

Any configuration change with this command automatically disables and enables the neighbors before the changes.

To configure this feature on Layer 3 VPNs, you must configure this feature in the context of the MPLS-enabled VR; this feature is not supported for BGP neighbors on the CE (VRF) side of the PE router.

Example

The following command applies the current input routing policy to the neighbors in the peer group outer:

```
configure bgp peer-group outer soft-reset in
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Support for Layer 3 VPNs was added in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure bgp peer-group source-interface

```
configure bgp peer-group peer-group-name source-interface [any |  
    ipaddress ipAddr]
```

Description

Configures the source interface for a peer group and all the neighbors of the peer group.

Syntax Description

<i>peer-group-name</i>	Specifies a peer group.
any	Specifies any source interface.
<i>ipAddr</i>	Specifies an interface.

Default

N/A.

Usage Guidelines

The source interface IP address must be a valid IP address of a VLAN configured on the switch.

Changes made to the parameters of a peer group are applied to all neighbors in the peer group.

This command applies to the current VR or VRF context.

After you enter this command, the switch automatically disables and enables the neighbors so that the changes can take effect.

Example

The following command configures the source interface for the peer group outer and its neighbors on 10.34.25.10:

```
configure bgp peer-group outer source-interface ipaddress 10.34.25.10
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the *BGP* feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure bgp peer-group timer

```
configure bgp peer-group peer-group-name 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

<i>peer-group-name</i>	Specifies a peer group.
keep-alive <i>seconds</i>	Specifies a keepalive time in seconds. Range is 0 to 21845.
hold-time <i>seconds</i>	Specifies a hold-time in seconds. Range is 0 and 3 to 65535.

Default

N/A.

Usage Guidelines

This command applies to the current VR or VRF context.

Changes made to the parameters of a peer group are applied to all neighbors in the peer group.

After you enter this command, the switch automatically disables and enables all neighbors in the peer group before the change takes effect.

Example

The following command configures the keepalive timer and hold timer values for the peer group outer and its neighbors:

```
configure bgp peer-group outer timer keep-alive 30 hold-time 90
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the *BGP* feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure bgp peer-group weight

```
configure bgp peer-group peer-group-name weight weight
```

Description

Configures the weight for the peer group and all the neighbors of the peer group.

Syntax Description

<i>peer-group-name</i>	Specifies a peer group.
<i>weight</i>	Specifies a <i>BGP</i> peer group weight. Range is 0 to 65,535.

Default

N/A.

Usage Guidelines

BGP selects routes based on the following precedence (from highest to lowest):

- higher weight
- higher local preference
- shortest length (shortest AS path)
- lowest origin code
- lowest MED
- route from external peer
- lowest cost to Next Hop
- lowest routerID

This command applies to the current VR or VRF context.

Example

The following command configures the weight for the peer group outer and its neighbors:

```
configure bgp peer-group outer weight 5
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure bgp restart address-family

```
configure bgp restart [add | delete] address-family [ipv4-unicast |  
ipv4-multicast | ipv6-unicast | ipv6-multicast]
```

Description

Configures the address family used with graceful [BGP](#) restart.

Syntax Description

add	Add the address family.
delete	Remove the address family.
address-family	Specifies an IPv4 or IPv6 unicast or multicast address family.

Default

The default is IPv4 unicast.

Usage Guidelines

Before you can enter this command, you must disable BGP services on the switch with the [disable bgp](#) command.

This command configures the address family participating in graceful BGP restart. An address family can be added or deleted. By adding an address family, BGP instructs the switch to preserve BGP routes of that address family during a graceful restart. The local OPEN message contains all the added address families.



Note

When graceful restart is enabled on the switch, the IPv4 unicast address family support is added by default. Graceful restart for other address families must be explicitly added using this command.

For BGP graceful restart to inter-operate with Cisco routers, any restarting routers connected to Cisco routers must be configured with the command, [enable bgp neighbor capability](#), in the following form: `enable bgp neighbor remoteaddr capability ipv4-unicast`. The command must be executed before BGP is enabled globally on the switch.

Example

The following command configures a router to add IPv4 unicast addresses to graceful BGP restarts:

```
configure bgp restart add address-family ipv4-unicast
```

History

This command was first available in ExtremeXOS 11.4.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure bgp restart restart-time

```
configure bgp restart restart-time seconds
```

Description

Configures the restart time used with graceful *BGP* restart. This is the maximum time a receiver router waits for a restarting router to come back up.

Syntax Description

<i>seconds</i>	Specifies the restart time. The range is 1 to 3600 seconds.
----------------	---

Default

The default is 120 seconds.

Usage Guidelines

Before you can enter this command, you must disable BGP services on the switch with the [disable bgp](#) command.

This command configures the restart timer. This timer is started on the receiver router when it detects the neighbor router is restarting (usually when the peer TCP session is reset). At that time, routes from the restarting router are marked as stale, but are preserved in the routing table. The timer is stopped when the restarting BGP neighbor goes to the ESTABLISHED state (it has finished restarting). If the timer expires, the stale routes are deleted.

Example

The following command configures the graceful BGP restart timer:

```
configure bgp restart restart-time 200
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure bgp restart stale-route-time

```
configure bgp restart stale-route-time seconds
```

Description

Configures the stale route timer used with graceful *BGP* restart. This is the maximum time to hold stale paths on receiver routers while its neighbor gracefully restarts.

Syntax Description

<i>seconds</i>	Specifies the stale route time. The range is 1 to 3600 seconds.
----------------	---

Default

The default is 360 seconds.

Usage Guidelines

Before you can enter this command, you must disable BGP services on the switch with the [disable bgp](#) command.

This command configures the stale route timer. This timer is started when the restarting BGP peer goes to the ESTABLISHED state after it restarts. The timer is stopped when the restarting BGP peer sends EOR messages for all address families. When the timer is stopped, or it expires, the stale routes are deleted.

Example

The following command configures the graceful BGP stale route timer:

```
configure bgp restart stale-route-time 400
```


History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure bgp restart update-delay

```
configure bgp restart update-delay seconds
```

Description

Configures the update delay timer used with graceful [BGP](#) restart. This is the maximum time to delay updating BGP routes to the local IP route table.

Syntax Description

<i>seconds</i>	Specifies the stale route time. The range is 1 to 3600 seconds.
----------------	---

Default

The default is 600 seconds.

Usage Guidelines

Before you can enter this command, you must disable BGP services on the switch with the [disable bgp](#) command.

This command configures the update delay timer. Usually, a restarting router waits to receive EOR messages from all the receiving BGP neighbors before it starts the route update. Otherwise, it does the route selection when the timer expires.

Example

The following command configures the graceful BGP update delay timer:

```
configure bgp restart update-delay 800
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure bgp restart

```
configure bgp restart [none | planned | unplanned | both | aware-only]
```

Description

Configures the router as a graceful *BGP* restart router.

Syntax Description

none	Do not act as a graceful BGP restart router.
planned	Only act as a graceful BGP restart router for planned restarts.
unplanned	Only act as a graceful BGP restart router for unplanned restarts.
both	Act as a graceful BGP restart router for both planned and unplanned restarts.
aware-only	Only act as a graceful BGP receiver (helper) router.

Default

The default is none; graceful restart is disabled.

Usage Guidelines

This command configures the router as a graceful BGP router. You can decide to configure a router to enter graceful restart for only planned restarts, for only unplanned restarts, or for both. Also, you can decide to configure a router to be a receiver only (which helps a restarting BGP router to perform the graceful restart process), and not to do graceful restarts itself.

After a graceful restart, the switch preserves the time stamps for all BGP routes in the RIB that were received before the stale timer expired. After restart, the capabilities for all BGP peers are renegotiated.



Note

End of Restart (EOR) messages are not sent to BGP peers if the graceful restart feature is disabled.

This command cannot be used while BGP is enabled globally on the switch.

Example

The following command configures a router to perform graceful BGP restarts only for planned restarts:

```
configure bgp restart planned
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure bgp routerid

```
configure bgp routerid router-identifier
```

Description

Changes the router identifier.

Syntax Description

<i>router identifier</i>	Specifies a router identifier in the IPv4 address format.
--------------------------	---

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):

- Higher weight
- Higher local preference
- Shortest length (shortest AS path)
- Lowest origin code
- Lowest MED
- Route from external peer
- Lowest cost to Next Hop
- Lowest router ID

This command applies to the current VR or VRF context.

Example

The following command changes the router ID:

```
configure bgp routerid 192.1.1.13
```



Note

To remove the configured bgp routerid, give routerid value as 0.0.0.0 i.e. configure bgp routerid 0.0.0.0.

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure bgp soft-reconfiguration

```
configure bgp soft-reconfiguration
```

Description

Immediately applies the route policy associated with the network command, aggregation, import, and redistribution.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command does not affect the switch configuration.

This command applies to the current VR or VRF context.

Example

The following command applies the route policy associated with the network command, aggregation, import, and redistribution:

```
configure bgp soft-reconfiguration
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the *BGP* feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure bootprelay

```
configure bootprelay [ {vlan [vlan_name]} [add ip_address | delete  
[ip_address] | all ]]
```

Description

This command configures DHCPv4 server/next hop relay for each *VLAN* IPv4 interfaces. This command is not applicable to IPv6 interfaces. Configuring bootprelay per VLAN v4 level is supported only on IPv4, and not on IPv6.

Syntax Description

bootprelay	BOOTP Relay service.
add	Adds <i>DHCP</i> BOOTP Relay server.
delete	Deletes DHCP BOOTP Relay server.
<i>ip_address</i>	IP address of bootp relay server.

Default

N/A.

Usage Guidelines

Use this command to configure the DHCPv4 server/next hop for each VLAN interface. The configuration applied to the VR level is populated to all VLAN v4 IPv4/v6 interfaces.

Example

The following example displays IPv6 bootprelay information:

```
# sh bootprelay configuration ipv4
DHCPv4 BOOTP Relay : Enabled on virtual router "VR-Default"
  Include Secondary      : Disabled
  BOOTP Relay Servers   : 10.127.6.243
  DHCP Relay Agent Information Option: Disabled
  DHCP Relay Agent Information Check : Disabled
  DHCP Relay Agent Information Policy: Replace

VLAN                                DHCPv4 BOOTP Relay
-----
VLAN "Default":
  BOOTP Relay                      : Enabled
VLAN "client":
  BOOTP Relay                      : Enabled
  BOOTP Relay Servers              : 10.1.1.1    10.127.6.101    10.127.6.243
  DHCP Relay Agent Information Option: Disabled
  DHCP Relay Agent Information Check : Disabled
  DHCP Relay Agent Information Policy: Replace
VLAN "client1":
  BOOTP Relay                      : Enabled
VLAN "dhcpv4server":
  BOOTP Relay                      : Enabled
VLAN "server":
  BOOTP Relay                      : Enabled
```

History

This command was first available in ExtremeXOS 15.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure bootprelay add

```
configure bootprelay {ipv4 | ipv6} add ip_address {vr vrid}
```

Description

Configures the addresses to which BOOTP requests should be directed.

Syntax Description

ipv4	DHCPv4 BOOTP Relay service (default).
ipv6	DHCPv6 BOOTP relay service.
<i>ip_address</i>	Specifies an IP address.
<i>vrid</i>	Specifies a VR or VRF name.

Default

If you do not specify a VR or VRF, the current VR context is used.

If you do not specify DHCPv4 or v6 BOOTP Relay service, DHCPv4 is used.

Usage Guidelines

After IP unicast routing has been configured, you can configure the switch to forward *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. Configure the addresses to which DHCP or BOOTP requests should be directed, using the following command: `configure bootprelay add ip_address`
3. Enable the DHCP or BOOTP relay function using the following command: `enable bootprelay`

Example

The following example configures BOOTP requests to be directed to 123.45.67.8:

```
configure bootprelay add 123.45.67.8
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure bootprelay delete

```
configure bootprelay {ipv4 | ipv6} delete [ip_address | all] {vr vrid}
```

Description

Removes one or all IP destination addresses for forwarding BOOTP packets.

Syntax Description

ipv4	DHCPv4 BOOTP Relay service (default).
ipv6	DHCPv6 BOOTP relay service.
<i>ip_address</i>	Specifies an IP address.
all	Specifies all IP address entries.
<i>vrid</i>	Specifies a VR or VRF name.

Default

If you do not specify a VR, the current VR context is used.

If you do not specify DHCPv4 or v6 BOOTP Relay service, DHCPv4 is used.

Usage Guidelines

None.

Example

The following command removes the destination address:

```
configure bootprelay delete 123.45.67.8
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure bootprelay dhcp-agent information check

```
configure bootprelay dhcp-agent information check
```

Description

Enables the *DHCP* relay agent option (option 82) checking.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

In some instances, a DHCP server may not properly handle a DHCP request packet containing a relay agent option. Use this command to prevent DHCP reply packets with invalid or missing relay agent options from being forwarded to the client.

To disable this check, use the following command:

```
unconfigure bootprelay dhcp-agent information check
```


Example

The following command configures the DHCP relay agent option check:

```
configure bootprelay dhcp-agent information check
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure bootprelay dhcp-agent information circuit-id port-information

```
configure bootprelay dhcp-agent information circuit-id port-information  
port_info port port
```

Description

Configures the circuit ID sub-option that identifies the port for an incoming *DHCP* request.

Syntax Description

<i>port_info</i>	Specifies a text string that becomes the circuit ID sub-option for the specified port. Specify a text string composed of 1 to 32 characters.
<i>port</i>	Specifies the port to which the circuit ID sub-option is assigned.

Default

The default *port_info* is encoded as ((slot_number * 1000) + port_number/portlindex). For example, if the DHCP request is received on port 3:12, the default circuit ID *port_info* value is 3012. On standalone switches, the slot number is one, so the default circuit ID *port_info* value is (1000 + port_number/portlindex). For example, the default *port_info* for port 3 on a standalone switch is 1003.

Usage Guidelines

The full circuit ID string uses the format *vlan_info-port_info* . To configure the *vlan_info* portion of the circuit ID string, use the following command:

```
configure bootprelay dhcp-agent information circuit-id vlan-information  
vlan_info {vlan} [vlan_name|all]
```

To display the *port_info* information, use the following command:

```
show bootprelay dhcp-agent information circuit-id port-information ports  
all
```

Example

The following command configures the circuit ID port_info value slot1port3 for port 1:3:

```
configure bootprelay dhcp-agent information circuit-id port-information slot1port3 port
1:3
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure bootprelay dhcp-agent information circuit-id vlan-information

```
configure bootprelay dhcp-agent information circuit-id vlan-information
    vlan_info {vlan} [vlan_name|all]
```

Description

Configures the circuit ID sub-option that identifies the VLAN for an incoming DHCP request.

Syntax Description

<i>vlan_info</i>	Specifies a text string that becomes the circuit ID sub-option for the specified VLAN. Specify a text string composed of 1 to 32 characters.
<i>vlan_name</i>	Specifies the VLAN to which the circuit ID sub-option is assigned.
all	Specifies that the <i>vlan_info</i> entered is to be used in the circuit ID sub-option for all VLANs.

Default

The default *vlan_info* for each VLAN is the VLAN ID or tag.

Usage Guidelines

The full circuit ID string uses the format *vlan_info-port_info* . To configure the port_info portion of the circuit ID string, use the following command:

```
configure bootprelay dhcp-agent information circuit-id port-information
port_info port port
```

To display the *vlan_info* information, use the following command:

```
show bootprelay dhcp-agent information circuit-id vlan-information
```

Example

The following command configures the circuit ID `vlan_info` value `VLANblue` for VLAN `blue`:

```
configure bootprelay dhcp-agent information circuit-id vlan-information VLANblue blue
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure bootprelay dhcp-agent information option

```
configure bootprelay dhcp-agent information option
```

Description

Enables the *DHCP* relay agent option (option 82).

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

After IP unicast routing has been configured, you can configure the switch to forward 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:

- Configure VLANs and IP unicast routing.
- Enable the DHCP or BOOTP relay function, using the following command: `enable bootprelay [{vlan} [vlan_name] | [{vr} vr_name] | all [{vr} vr_name]}`
- Configure the addresses to which DHCP or BOOTP requests should be directed, using the following command: `configure bootprelay add ip_address {vr vrid}`

Configure the DHCP relay agent option (option 82), using the following command:

```
configure bootprelay dhcp-agent information option
```

To disable the DHCP relay agent option (option 82), use the following command: `unconfigure bootprelay dhcp-agent information option`

Example

The following example configures the DHCP relay agent option:

```
configure bootrelay dhcp-agent information option
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure bootrelay dhcp-agent information policy

```
configure bootrelay dhcp-agent information policy [drop | keep |  
replace]
```

Description

Configures the *DHCP* relay agent option (option 82) policy.

Syntax Description

drop	Specifies to drop the packet.
keep	Specifies to keep the existing option 82 information in place.
replace	Specifies to replace the existing data with the switch's own data.

Default

Replace.

Usage Guidelines

Use this command to set a policy for the relay agent. Packets can be dropped, the option 82 information can be replaced (the default), or the packet can be forwarded with the information unchanged.

Example

The following command configures the DHCP relay agent option 82 policy to keep:

```
configure bootrelay dhcp-agent information policy keep
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure bootrelay dhcp-agent information remote-id

```
configure bootrelay dhcp-agent information remote-id [remote_id |
system-name] {vr vrid}
```

Description

Configures the remote ID sub-option that identifies the relaying switch for *DHCP* requests and replies.

Syntax Description

<i>remote_id</i>	Specifies a text string that becomes the remote ID sub-option for the switch. Specify a text string composed of 1 to 32 characters.
system-name	Specifies that the switch name is used as the remote ID sub-option for the switch.
<i>vrid</i>	Specifies the VR on which to configure the remote ID sub-option.

Default

The switch MAC address.

Usage Guidelines

To display the remote-ID, use the following command: `show bootrelay`

Example

The following example configures the remote ID sub-option to specify the switch name in DHCP requests and replies:

```
configure bootrelay dhcp-agent information remote-id system-name
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure bootrelay dhcp-agent source-vlan

```
configure bootrelay dhcp-agent source-vlan {vlan_name} {vr vrid}
```

Description

Configures a source VLAN to use as the source IP address in the BOOTPrelay packet.

Syntax Description

bootprelay	Specifies BOOTPrelay agent information option.
dhcp-agent	Specifies DHCP agent.
source-vlan	Specifies using the IP address of the configured loopback VLAN as the gateway IP (giaddr) field when BOOTPrelay is used for the anycast VLAN.
<i>vlan_name</i>	Specifies the loopback VLAN name.
vr	Specifies using a specific virtual router ID.
<i>vrid</i>	Specifies the virtual router ID.

Default

N/A.

Usage Guidelines

When a VLAN is configured with an anycast MAC address and has an anycast IP address, the source address must be a unique reachable IP address. Without a unique IP address, the DHCP reply from the server might not reach the DHCP relay ExtremeXOS device. The source VLAN must be a loopback VLAN that is configured on the specified VRF. If a source VLAN is not configured, then a loopback VLAN is automatically picked for anycast VLANs receiving a DHCP request.

To view the selected source VLAN, use the command `show bootprelay configuration {ipv4 | ipv6} {{vlan vlan_name } | {vr vr_name}}` .

Example

The following example configures the VLAN "vlan1" to use as the source IP address in the BOOTPrelay packet:

```
# configure bootprelay dhcp-agent source-vlan vlan1
```

History

This command was first available in ExtremeXOS 30.7.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure bootprelay include-secondary

```
configure bootprelay {ipv4 | ipv6} include-secondary {sequential |  
  parallel | off} {vr vr_name}
```

Description

Configures DHCP smart relay mode, and includes a secondary IP address as the giaddr at the VR level.

Syntax Description

ipv4	Specifies DHCPv4 BOOTP Relay service (default).
ipv6	Specifies DHCPv6 BOOTP Relay service.
include-secondary	(Optional) Uses both primary and secondary address(es) of the client <u>VLAN</u> as gateway address.
sequential	Uses primary and secondary address(es) of client VLAN in sequence after 3 retries (default if include-secondary is on).
parallel	Uses primary and secondary address(es) of client VLAN in parallel.
off	Disables use of both primary and secondary address(es) of client VLAN as gateway address (default).
vr	Specifies a virtual router ID.
<i>vr_name</i>	Specifies the virtual router.

Default

IPv4 is the default relay service.

The default value is **off**, but **sequential** is the default if **include-secondary** is on.

Usage Guidelines

Use this command to configure DHCP smart relay mode, and to include a secondary IP address as giaddr at the VR level.

Example

The following example configures DHCPv4 BOOTP Relay service to use both primary and secondary addresses of the client VLAN as the gateway address. By default, the command specifies that you use the primary and secondary addresses of the client VLAN in sequence after three retries.

```
configure bootprelay ipv4 include-secondary sequential
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure bootprelay ipv6 option interface-id

```
configure bootprelay ipv6 option [ interface-id ] [identifier_string |
system_name | none] [vlan vlan_name | all]
```

Description

This command configures the option **interface-id** as described in RFC-4649 to an IPv6 bootp relay/*DHCP* relay agent.

Syntax Description

bootprelay	BOOTP Relay Service
ipv6	DHCPv6 BOOTP Relay Service
option	DHCPv6 BOOTP Relay options
interface-id	Interface identifier option.
<i>interface_id_string</i>	Interface identifier string.
none	Identifier defaults to 802.1Q <i>VLAN</i> ID.
all	All VLANs.

Default

802.1Q VLAN ID if not configured.

Usage Guidelines

Use this command to configure the option **interface-id** as described in RFC-4649 to an IPv6 BOOTP relay/DHCP relay agent. After receiving an IPv6 BOOTP/DHCP request packet on the specified VLAN, the agent adds the configured identifier to the packet and passes it to the server. If this option is configured to be as **system-name**, the switch name is used as the remote-id. The same can be unconfigured using the **none** option. After unconfiguring this option, the switch MAC address (the default value) is used as remote-id. This option can be configured or unconfigured to a specified VLAN or to all VLANs.

Example

```
* Switch # show bootprelay ipv6
BOOTP Relay: DHCPv6 BOOTP Relay enabled on virtual router "VR-Default"
  BOOTP Relay Servers : 2001:0db8:85a3:0000:0000:8a2e:0370:7334
                        2001:0db8:85a3:0000:0000:8a2e:0370:7335
                        2001:0db8:85a3:0000:0000:8a2e:0370:7336
                        2001:0db8:85a3:0000:0000:8a2e:0370:7337
VLAN "Default":
  BOOTP Relay      : Disabled
```



```

VLAN "v1":
  BOOTP Relay      : Enabled
  Interface ID     : v1-12
  Remote ID        : v1_remId
VLAN "v2":
  BOOTP Relay      : Enabled
  Interface ID     : 100 (Default)
  Remote ID        : 00:04:96:52:A7:1B (Default)

```

History

This command was first available in ExtremeXOS 15.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure bootprelay ipv6 option remote-id

```

configure bootprelay ipv6 option [remote-id] [identifier_string] |
  system-name | none] [vlan vlan_name | vlan all]

```

Description

This command configures the **remote-id** option as described in RFC-4649 to an IPv6 BOOTP relay/
DHCP relay agent.

Syntax Description

bootprelay	BOOTP Relay Service
ipv6	DHCPv6 BOOTP Relay Service
option	DHCPv6 BOOTP Relay options
remote-id	Remote-ID sub-option to identify remote host
<i>remote_id_string</i>	Remote ID String
<i>system-name</i>	System Name string
none	Identifier defaults System MAC address
all	All <i>VLANs</i>

Default

System MAC address if not configured.

Usage Guidelines

Use this command to configure the **remote-id** option as described in RFC-4649 to an IPv6 BOOTP relay/DHCP relay agent. After receiving an IPv6 BOOTP/DHCP request packet on the specified VLAN, the agent adds the configured identifier to the packet and passes it to the server. If this option is

configured to be as **system-name**, the switch name is used as the remote-id. The same can be unconfigured using the **none** option. After unconfiguring this option, the switch MAC address (the default value), is used as remote-id. This option can be configured or unconfigured to a specified VLAN or to all VLANs.

Example

```
* Switch # show bootprelay ipv6
BOOTP Relay: DHCPv6 BOOTP Relay enabled on virtual router "VR-Default"
  BOOTP Relay Servers : 2001:0db8:85a3:0000:0000:8a2e:0370:7334
                        2001:0db8:85a3:0000:0000:8a2e:0370:7335
                        2001:0db8:85a3:0000:0000:8a2e:0370:7336
                        2001:0db8:85a3:0000:0000:8a2e:0370:7337

VLAN "Default":
  BOOTP Relay      : Disabled
VLAN "v1":
  BOOTP Relay      : Enabled
  Interface ID     : v1-12
  Remote ID        : v1_remId
VLAN "v2":
  BOOTP Relay      : Enabled
  Interface ID     : 100 (Default)
  Remote ID        : 00:04:96:52:A7:1B (Default)
```

History

This command was first available in ExtremeXOS 15.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure bootprelay ipv6 prefix-delegation snooping add

```
configure bootprelay ipv6 prefix-delegation snooping add ipv6_prefix
ipv6Gateway {vlan} vlan_name valid-time valid_time
```

Description

Adds information about a snooped IPv6 delegated prefix on a VLAN.

Syntax Description

<i>ipv6_prefix</i>	Specifies the IPv6 prefix (/prefix length) to be added.
<i>ipv6Gateway</i>	Specifies the IPv6 gateway address.
<i>vlan_name</i>	Specifies the VLAN.
<i>valid_time</i>	Time, in seconds, that the delegated IPv6 prefix is valid.

Default

N/A

Usage Guidelines

Allows you to add a particular IPv6 delegated prefix to snoop if the prefix was issued or renewed during reboot. If the prefix has been snooped earlier, this command renews the valid time for the prefix.

To set the specified prefix to always be valid, set the **valid-time** parameter to 0.

Before adding an IPv6 delegated prefix to snoop, you must enable IPv6 BOOTP relay and prefix snooping using `enable bootprelay ipv6` and `configure bootprelay ipv6 prefix-delegation snooping` .

Example

The following example adds prefix /56.

```
configure bootprelay ipv6 prefix-delegation snooping add 5001:db8:3553:bf00::/56
fe80::a440:cf5:c05b:d324 vlan vl valid-time 300
```

History

This command was first available in ExtremeXOS 15.7.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure bootprelay ipv6 prefix-delegation snooping

```
configure bootprelay ipv6 prefix-delegation snooping [on {vlan}vlan_name
| off [{vlan}vlan_name | vlan all] ]
```

Description

Enables and disables snooping of IPv6 prefixes delegated via *DHCP*.

Syntax Description

on	Enables snooping of IPv6 prefixes delegated via DHCP on the specified <i>VLAN</i> .
<i>vlan_name</i>	Specifies the VLAN.
off	Disables snooping of IPv6 prefixes delegated via DHCP on the specified VLAN.
vlan all	Disables snooping of IPv6 prefixes delegated via DHCP on all VLANs.

Default

By default, snooping of IPv6 prefixes is off.

Usage Guidelines

You can enable snooping on a specific VLAN.

You can disable the snooping on a specific VLAN or all VLANs.

Example

The following example disables snooping of IPv6 prefixes on all VLANs.

```
configure bootprelay ipv6 prefix-delegation snooping off vlan all
```

History

This command was first available in ExtremeXOS 15.7.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure bootprelay vlan include-secondary

```
configure bootprelay {ipv4 | ipv6} {vlan vlan_name} include-secondary
{sequential | parallel | off}
```

Description

Configures DHCP smart relay mode to include secondary IP address as giaddr at VLAN level.

Syntax Description

ipv4	Specifies DHCPv4 BOOTP Relay service (default).
ipv6	Specifies DHCPv6 BOOTP Relay service.
vlan	Configure BOOTP relay for this VLAN, and overrides the VR level configuration.
<i>vlan_name</i>	Specifies the VLAN name.
include-secondary	(Optional) Use both primary and secondary address(es) of the client VLAN as gateway address.
sequential	Use primary and secondary address(es) of client VLAN in sequence after 3 retries (default if include-secondary is on).
parallel	Use primary and secondary address(es) of client VLAN in parallel.
off	Disable use of both primary and secondary address(es) of client VLAN as gateway address (default).

Default

IPv4 is the default relay service.

off is the default value, but **sequential** is the default if **include-secondary** is on.

Usage Guidelines

Use this command to configure DHCP smart relay mode to include the secondary IP address as giaddr at the VLAN level.

Example

The following command configures DHCPv4 BOOTP Relay service for the "vlan_100" VLAN, and uses both primary and secondary address(es) of the client VLAN as gateway address. This overrides the VR level configuration.

```
configure bootprelay ipv4 vlan vlan_100 include-secondary
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure cdp cos-extend ports

```
configure cdp cos-extend cos_value ports [port_list | all]
```

Description

This command configures COS extended support on the IP phone. This information will be sent to the IP phone from the ExtremeXOS switch by trust TLV and COS TLV.

Syntax Description

cos_value	COS value range from 0 to 7.
port_list	Port list separated by a comma or "-";

Default

0.

Usage Guidelines

None.

Example

The following example sets the COS TLV value as 4 for port 5 in the ExtremeXOS switch, which will be used by the IP phone to override priority received from PC or the attached device.

```
configure cdp cos-extend 4 ports 5
```

The following example sets the COS TLV value to default for port 5 in the ExtremeXOS switch.

```
configure cdp cos-extend 0 ports 5
```

History

This command was first available in ExtremeXOS 21.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure cdp device-id

```
configure cdp device-id [device_id | system-mac | system-name]
```

Description

Configures the device ID only in CDP.

Syntax Description

<i>device-id</i>	Unique device identifier to be used in CDP.
system-mac	Use system MAC address as the device identifier.
system-name	Use sysName as the device identifier (default).

Default

system-name.

Usage Guidelines

Use this command to configure the Device ID. If you do not configure it, the MAC address is used as the Device ID. This configuration of device ID is only used in the CDP .

Example

The following command configures the device ID as the MAC address:

```
configure cdp device-id system-mac
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure cdp frequency

```
configure cdp frequency seconds
```

Description

Enables CDP on a port.

Syntax Description

<i>seconds</i>	Specifies the transmit frequency in seconds. The range is 5,254 seconds. The default value is 60 seconds.
----------------	---

Default

60 seconds.

Usage Guidelines

Example

The following command configures the CDP frequency as two minutes:

```
configure cdp frequency 120
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure cdp hold-time

```
configure cdp hold-time seconds
```

Description

Configures the hold time of the neighbor information .

Syntax Description

<i>seconds</i>	Duration in seconds that receiver must keep this packet. The range is 10-255 and the default is 180 seconds.
----------------	--

Default

60 seconds.

Usage Guidelines

Use this command to configure the hold time of the neighbor information for which a receiving device should hold information before discarding it.

Example

The following command configures the CDP hold time as two minutes:

```
configure cdp hold-time 120
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure cdp management-address

```
configure cdp management-address [{vlan} vlan_name | vlan vlan_id]  
    {primary-ip | secondary-ip secondary_ip_address}
```

Description

Configures a specified VLAN's IP address as the management address to be advertised by Cisco Discovery Protocol (CDP).

Syntax Description

vlan	Specifies a VLAN for the management IP address.
<i>vlan_name</i>	Specifies a VLAN name for the management IP address (default is "Mmgt").
<i>vlan_id</i>	Specifies a VLAN ID for the management IP address.
primary-ip	CDP advertises the primary IP address of the specified VLAN (default). The specified VLAN must be already configured with at least one primary IPv4 or IPv6 address.
secondary-ip	Specifies that CDP advertises the secondary IP address of the specified VLAN. The specified secondary IP address must already be configured on the specified VLAN.
<i>secondary_ip_address</i>	Specifies the secondary IP address of the specified VLAN.

Default

By default, the Management VLAN's IP address is advertised by CDP.

If you do not specify, CDP advertises the primary IP address of the specified VLAN.

Usage Guidelines

If the Management VLAN IP address is not configured, you can specify any user-defined VLAN's IP address or front panel port VLAN's IP address as the management address for the CDP protocol.

This command dictates the management address to be advertised by the CDP protocol; the equivalent command for *LLDP* is [configure lldp management-address](#) on page 761.

To use this command, the specified VLAN must already exist. The management IP address configuration is removed if the specified VLAN is deleted, or if the primary IP address of the specified VLAN is deleted (if **primary-ip** configured), or if the specified secondary IP address of the specified VLAN is deleted (if **secondary-ip** configured).

If **primary-ip** is configured and the specified VLAN has multiple primary IP addresses (IPv4 and IPv6), then CDP advertises the first primary IP address that exists in the address table. If IPv4 is not configured, CDP advertises the first IPv6 address.

If **secondary-ip** is configured and the specified VLAN has multiple secondary IP addresses, then CDP advertises only the specified secondary IP address of the configuration.

Example

The following example configures the primary IP address of the VLAN "vlan1" as the management address to be advertised by CDP protocol:

```
configure cdp management-address vlan vlan1 primary-ip
```

History

This command was first available in ExtremeXOS 22.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure cdp power-available ports

```
configure cdp power-available [advertise | no-advertise] ports  
    [port_list | all]
```

Description

This command configures the advertising status of the power available TLV on CDP ports.

Syntax Description

advertise	Specifies to send the TLV to neighbors.
<i>no-advertise</i>	Specifies not to send the TLV to neighbors.
<i>port_list</i>	Port list separated by a comma or - .

Default

No-advertise.

Usage Guidelines

This command is for *PoE* switches.

Example

The following example advertises the Power Available TLV on port 1:

```
configure cdp power-available advertise ports 1
```

The following example does not advertise the Power Available TLV on port 1:

```
configure cdp power-available no-advertise ports 1
```

History

This command was first available in ExtremeXOS 21.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure cdp trust-extend ports

```
configure cdp trust-extend [untrusted | trusted] ports [port_list | all]
```

Description

This command configures trust mode support for the IP phone. This information will be sent to the IP phone from the ExtremeXOS switch by trust TLV.

Syntax Description

untrusted	Instructs attached IP phone to overwrite priority received from PC with configured COS value.
trusted	Instructs IP phone to trust the priority received from PC or the attached device.
<i>port_list</i>	Port list separated by a comma or "-";

Default

Trusted.

Usage Guidelines

None.

Example

The following example sets the trust TLV value as trusted for port 5 in the ExtremeXOS switch, which will be used by the IP phone to not change the priority received from the PC or attached device.

```
configure cdp trust-extend trusted ports 5
```

The following example sets the trust mode to untrusted for port 5 in the ExtremeXOS switch, which will be used by the IP phone to override priority received from the PC or attached device.

```
configure cdp trust-extend untrusted ports 5
```

History

This command was first available in ExtremeXOS 21.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure cdp voip-vlan ports

```
configure cdp voip-vlan advertise [solicited | unsolicited] [vlan_name |  
vlan_id | dot1p | untagged | none] ports [port_list | all]
```

Description

This command configures voice *VLAN*, for voice traffic from the IP phone in one or more ports in ExtremeXOS switch. This information will be sent to IP phone from the ExtremeXOS switch by VOIP Reply TLV.

Syntax Description

<i>vlan_name</i>	VLAN name.
<i>vlan_id</i>	VLAN ID tag between 1 and 4,094.
advertise	Configures when TLVs are sent to neighbors.
solicited	Send TLVs to neighbors only when requested (default).
unsolicited	Send TLVs to neighbors without waiting for a request.
dot1p	Instructs IP phone to send dot1p tagged voice traffic.
untagged	Instructs IP phone to send untagged voice traffic.
none	No VLAN information is sent in CDP PDUs.
<i>port_list</i>	Port list is separated by a comma or "-";
all	Instructs IP phone to send all traffic.

Default

By default, voice VLAN reply TLV are sent to neighbors only when requested.

Usage Guidelines

None.

Example

The following example sets the VOIP VLAN reply TLV value value as default in ExtremeXOS for port, which will be used by the IP phone for voice traffic:

```
Configure cdp voip-vlan "Default" ports 5
```

The following example sets the VOIP VLAN reply TLV value as priority tagged in ExtremeXOS switch for port 5, which will be used by the IP phone for voice traffic.

```
configure cdp voip-vlan dot1p ports 5
```

The following example sets the VOIP VLAN reply TLV value as untagged in ExtremeXOS switch, which will be used by the IP phone for voice traffic.

```
configure cdp voip-vlan untagged ports 5
```

The following example sets the VOIP VLAN reply TLV value as none in ExtremeXOS switch, this will not transmit any VLAN information TLV to the IP phone.

```
configure cdp voip-vlan none ports 5
```

The following example sets the VOIP VLAN reply TLV value as VLAN Id 1 in ExtremeXOS switch for port 5, which will be used by the IP phone for voice traffic.

```
configure cdp voip-vlan 1 ports 5
```

To configure VoIP VLAN as unsolicited on port 10:

```
configure cdp voip-vlan advertise unsolicited ports 10
```

History

This command was first available in ExtremeXOS 21.1.

Ability to send voice VLAN reply TLV without receiving voice VLAN request TLV (unsolicited) was added in ExtremeXOS 22.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure cfm domain add association integer

```
configure cfm domain domain_name add association integer int [vlan
vlan_name|vman vman_name]
```

Description

Creates a maintenance association (MA) related to a specified maintenance domain (MD). This command supports the 2-octet integer MA format.

Syntax Description

<i>domain_name</i>	Specifies the domain you want to associate with this MA.
<i>int</i>	Enter an integer to name the MA. The range is 0 to 65535.
<i>vlan_name</i>	Specifies the <u>VLAN</u> you want to assign to this MA. Each MA contains only one VLAN, VMAN, BVLAN or SVLAN.
<i>vman_name</i>	Specifies the VMAN you want to assign to this MA.

Default

N/A.

Usage Guidelines

All ports configured on the specified VLAN are now CFM ports in the specified MA.

You add the MA, or association, to the domain, and the MA uses the MD level assigned to the domain. Each MA can belong to only one domain, but several MAs can belong to a given domain. The MA is unique within a given domain.

Example

The following command creates a 2-octet integer MA (350) that associates the domain brazil and the VLAN admin:

```
configure cfm domain brazil add association integer 350 vlan admin
```

History

This command was first available in ExtremeXOS 11.4.

The SVLAN option was added in ExtremeXOS 12.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure cfm domain add association meg

```
configure cfm domain domain_name add association [meg meg_name]
```

Description

Creates a maintenance association (MA) related to a specified maintenance domain (MD). This command supports the MEG MA format.

Syntax Description

meg	ITU-T Y.1731 Maintenance Entity Group.
<i>meg_name</i>	MEG name, maximum of 12 characters with 6 bytes ITU Carrier Code and 6 bytes Organization specific Unique MEG ID Code.

Default

N/A.

Usage Guidelines

All ports configured on the specified MEG are now CFM ports in the specified MA. You add the MA, or association, to the domain, and the MA uses the MD level assigned to the domain. Each MA can belong to only one domain, but several MAs can belong to a given domain. The MA is unique within a given domain.

History

This command was first available in ExtremeXOS 15.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure cfm domain add association string

```
configure cfm domain domain_name add association string name [vlan
vlan_name | vman vman_name]
```

Description

Creates a maintenance association (MA) related to a specified maintenance domain (MD). This command supports the character string MA format.

Syntax Description

<i>domain_name</i>	Specifies the domain you want to associate with this MA.
string	Enter up to 45 alphanumeric characters to name the MA.
<i>vlan_name</i>	Specifies the <i>VLAN</i> you want to assign to this MA. Each MA contains only one VLAN, VMAN, or BVLAN.
<i>vman_name</i>	Specifies the VMAN you want to assign to this MA.

Default

N/A.

Usage Guidelines

All ports configured on the specified VLAN are now CFM ports in the specified MA.

You add the MA, or association, to the domain, and the MA uses the MD level assigned to the domain. Each MA can belong to only one domain, but several MAs can belong to a given domain. The MA is unique within a given domain.

Example

The following command creates an MA named service that associates the MD spain and the VLAN finance:

```
configure cfm domain service add association string spain vlan finance
```

History

This command was first available in ExtremeXOS 11.4.

The SVLAN option was added in ExtremeXOS 12.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure cfm domain add association vlan-id

```
configure cfm domain domain_name add association vlan-id vlanid [vlan
vlan_name|vman vman_name]
```

Description

Creates a maintenance association (MA) related to a specified maintenance domain (MD). This command supports the VLAN ID MA format.

Syntax Description

<i>domain_name</i>	Specifies the domain you want to associate with this MA.
<i>vlanid</i>	Specifies the VLAN ID.
<i>vlan_name</i>	Specifies the VLAN you want to assign to this MA. Each MA contains only one VLAN, VMAN, or BVLAN.
<i>vman_name</i>	Specifies the VMAN you want to assign to this MA.

Default

N/A.

Usage Guidelines

All ports configured on the specified VLAN are now CFM ports in the specified MA.

You add the MA, or association, to the domain, and the MA uses the MD level assigned to the domain. Each MA can belong to only one domain, but several MAs can belong to a given domain. The MA is unique within a given domain.

History

This command was first available in ExtremeXOS 12.1.

The SVLAN option was added in ExtremeXOS 12.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure cfm domain add association vpn-id oui index

```
configure cfm domain domain_name add association vpn-id oui oui index
index [vlan vlan_name | meg meg_name | vman vman_name]
```

Description

Creates a maintenance association (MA) related to a specified maintenance domain (MD). This command supports the RFC 2685 VPN ID MA format.

Syntax Description

<i>domain_name</i>	Specifies the domain you want to associate with this MA.
association	IEEE 802.1ag Maintenance Association or ITU-T Y.1731 Maintenance Entity Group
<i>oui</i>	Enter a virtual private network (VPN) Organizational Unique Identifier (OUI) in the format XX:XX:XX as part of the name for the MA.
<i>index</i>	Enter the 32-bit VPN index you want to append to the OUI to name the MA. The range is 0 to 4294967295.
<i>vlan_name</i>	Specifies the <i>VLAN</i> you want to assign to this MA. Each MA contains only one VLAN, VMAN, or BVLAN.
<i>vman_name</i>	Specifies the VMAN you want to assign to this MA.
meg	ITU-T Y.1731 Maintenance Entity Group.
<i>meg_name</i>	MEG name, maximum of 12 characters with 6 bytes ITU Carrier Code and 6 bytes organization specific unique MEG ID code.

Default

N/A.

Usage Guidelines

All ports configured on the specified VLAN are now CFM ports in the specified MA. You add the MA, or association, to the domain, and the MA uses the MD level assigned to the domain. Each MA can belong to only one domain, but several MAs can belong to a given domain. The MA is unique within a given domain.

Example

The following command creates an MA with the VPN ID of 11:22:33 50 that associates the domain spain and the VLAN accounting:

```
configure cfm domain spain add association vpn-id oui 11:22:33 index 50 vlan accounting
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure cfm domain association add remote-mep

```
configure cfm domain domain_name association association_name add
remote-mep mepid { mac_address mac_address }
```

Description

Allows you to add a remote MEP with the given MEP ID and MAC address to an existing association.

Syntax Description

<i>domain_name</i>	Enter the domain associated with the MA you are configuring.
<i>association_name</i>	IEEE 802.1ag or ITU-T Y.1731 association name.
<i>mepid</i>	Enter the MEP ID of the remote MEP being added. The range is 1 to 8191.
<i>mac_address</i>	Specifies the MAC address for the remote MEP being added.

Default

N/A.

Usage Guidelines

Use this command to add a remote MEP with given MEP ID and MAC address to an existing association. Use the `show cfm detail` command to verify your configuration.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure cfm domain association add

```
configure cfm domain domain_name association association_name [ports
port_list add [[end-point [up|down] mepid { group group_name } ] ] |
[intermediate-point]]
```

Description

This command allows you to create an up MEP, down MEP, intermediate-point (MIP) on a maintenance association, a group. You can also combine different maintenance points.

Combining different Maintenance points is restricted per the following:

- Up MEP and Down MEP in a single association is not allowed.
- Down MEP and MIP in a single association is not allowed.
- More than one Up MEP in a single association is not allowed.
- Up MEP and MIP in a single association is allowed.
- More than one Down MEP in a single association is allowed.
- A group can be created while creating a MEP.
- With CFM Support over VPLS, this command is used to associate pseudo wires of a VPLS service instance to an association & domain.
- Portlist can have only one port configured for a MEP configuration but can have multiple ports in MIP configuration, when Hwaoam is supported on the system.

Syntax Description

<i>domain_name</i>	Enter the domain associated with the MA you are configuring.
<i>association_name</i>	IEEE 802.1ag or ITU-T Y.1731 association name.
<i>port_list</i>	Specifies the port number(s).
up	Enter the port to be the UP port of the MA; this MEP sends CCM messages to all ports—other than the sending switch port—in this MA on this switch.
down	Enter the port to be the DOWN port of the MA; this MEP sends CCM messages out of the configured physical port.
<i>mepid</i>	Specifies a value for this MEP. The range is 1 to 8191. NOTE: On each MA, each MEPID must be unique.
group	CFM group that binds an LMEP to RMEPs. If not specified, the client does not receive events from the respective RMEPs.
<i>group_name</i>	Group name, maximum of 31 characters.

Default

N/A.

Usage Guidelines

These ports must already be in the MA (VLAN or VMAN) prior to assigning a MEP function to them. If you try to assign a port not in the MA as an end-point, the system returns the following message:

The following port(s) <portlist> are not part of the associations VLAN.



Note

Ensure that you assigned the port number correctly to the UP MEP and to the DOWN MEP, or the CCM messages go in the wrong direction.

Each MA needs at least two MEPs that can reach each other to exchange CCM messages.

You can also combine different maintenance points. The following are CLI restrictions on MP combinations:

- DOWN and UP MEP cannot be present on the same association
- DOWN MEP and MIP cannot be present on the same association
- UP MEP and MIP can be present on the same association
- Only one UP MEP is allowed in an association
- Multiple DOWN MEPs are allowed in an association

You can configure a total of 32 MIPs on a single switch.

Use the `show cfm` command to verify your configuration.

Example

The following command configures port 1:20 as a MIP on the 350 association in the spain domain:

```
configure cfm domain spain association 350 ports 1:20 add intermediate-point
```

The following command configures port 5:10 to be the UP MEP on the test association in the brazil domain, with a mepid of 500:

```
configure cfm domain brazil association test ports 5:10 add end-point up 500
```

History

This command was first available in ExtremeXOS 11.4.

This command was updated in ExtremeXOS 15.2 to include the optional group parameter.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure cfm domain association delete remote-mep

```
configure cfm domain domain_name association association_name delete
remote-mep mepid
```

Description

Allows you to delete a remote MEP for a specific MEP ID and MAC address.

Syntax Description

<i>domain_name</i>	Enter the domain associated with the MA you are configuring.
<i>association_name</i>	IEEE 802.1ag or ITU-T Y.1731 association name.
<i>mepid</i>	Enter the MEP ID of the remote MEP that is to be deleted.

Default

N/A.

Usage Guidelines

Use this command to delete a remote MEP of an MA for a specific MEP ID.

Use the `show cfm detail` command to verify your configuration.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure cfm domain association delete

```
configure cfm domain domain_name association association_name [ports
    port_list delete [[end-point [up|down]] | [intermediate-point] ] ]
```

Description

Deletes a maintenance end point (MEP) or maintenance intermediate point (MIP) from that MA.

Syntax Description

<i>domain_name</i>	Enter the domain associated with the MA you are configuring.
<i>association_name</i>	IEEE 802.1ag or ITU-T Y.1731 association name
<i>port_list</i>	Specifies the port number(s).
up	Specifies that an UP MEP is to be deleted.
down	Specifies that a DOWN MEP is to be deleted.

Default

N/A.

Usage Guidelines

Use this command to delete an MEP or MIP.

If the VPLS option is chosen then the CFM deletes all the VPLS-based MIPs.

Use the [show cfm](#) command to verify your configuration.

Example

The following command deletes port 5:12 as an MIP on the test association in the brazil domain:

```
configure cfm domain brazil association test ports 5:12 delete intermediate-point
```

The following command deletes an UP MEP on port 5:10 on the test association in the brazil domain:

```
configure cfm domain brazil association test ports 5:10 delete end-point up
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure cfm domain association destination-mac-type

```
configure cfm domain domain-name association association_name
destination-mac-type [unicast | multicast]
```

Description

Allows you to choose the destination MAC type for sending CFM PDUs for an MA.

Syntax Description

<i>domain_name</i>	Enter the domain associated with the MA you are configuring.
<i>association_name</i>	IEEE 802.1ag or ITU-T Y.1731 association name.
unicast	CFM PDUs are sent to the unicast MAC address configured in static remote MEP creation.
multicast	CFM PDUs are sent to the standard multicast destination address.

Default

Multicast.

Usage Guidelines

Use this command to change the MAC type on a previously configured MA. If multicast is selected, CFM PDUs are sent to the standard multicast destination. If unicast is selected, CFM PDUs are sent to the unicast MAC address configured in static remote MEP creation.

Use the `show cfm` command to verify your configuration.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure cfm domain association end-point add group

```
configure cfm domain domain-name association association-name ports
port-list end-point [up | down] add group group_name
```

Description

This command allows you to create a group for an existing local end-point.

Syntax Description

<i>domain_name</i>	Enter the domain associated with the MA you are configuring.
<i>association_name</i>	IEEE 802.1ag or ITU-T Y.1731 association name.
<i>port_list</i>	Enter the port number you want to configure as either an UP or DOWN MEP.

Default

N/A.

Usage Guidelines

Use this command to add a group to the association.

Example

```
configure cfm domain "MD1" association "MD1v1" ports 17 end-point down add group
"eapsCfmGrp"
```

History

This command was first available in ExtremeXOS 15.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure cfm domain association end-point delete group

```
configure cfm domain domain_name association association_name ports
port_list end-point [up|down] delete group [group_name | all ]
```

Description

This command allows you to delete one or all groups.

Syntax Description

<i>domain_name</i>	Enter the domain associated with the MA you are configuring.
<i>association_name</i>	IEEE 802.1ag or ITU-T Y.1731 association name.
<i>port_list</i>	Enter the port number you want to configure as either an UP or DOWN MEP.
delete	Delete configuration from the association

Default

N/A.

Usage Guidelines

Use this command to delete one or all groups from the association.

Example

```
configure cfm domain "MD1" association "MD1v1" ports 17 end-point down delete group
"eapsCfmGrp"
```

History

This command was first available in ExtremeXOS 15.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure cfm domain association end-point transmit-interval

```
configure cfm domain domain_name association association_name {ports
    port_list end-point [up | down]} transmit-interval [3|10|100|1000|
    10000|60000|600000]
```

Description

Allows you to change time interval for an MEP to send out a CCM. We recommend configuring this value as at least 1 second.

Syntax Description

<i>domain_name</i>	Enter the domain associated with the MA you are configuring.
<i>association_name</i>	IEEE 802.1ag or ITU-T Y.1731 association name.
<i>port_list</i>	Enter the port number of the MEP on which you are changing the time interval it sends out a CCM.
up	Enter this variable if you are changing the time interval for sending a CCM on an UP MEP.
down	Enter this variable if you are changing the time interval for sending a CCM on a DOWN MEP.

Default

1000 ms.

Usage Guidelines

Use this command to change the time interval between sending out CCMs on a previously configured UP or DOWN MEP. If you attempt to change the interval on a port that is either not an MEP or having wrong MEP type, the system returns an error message.



Note

We recommend that you use a transmit interval of at least 1 second (1000 ms).

The receiving system also uses this value multiplied by 3.5 to determine when the MEP is no longer alive.

Use the `show cfm` command to verify your configuration and the `show cfm detail` command to display the configured lifetime.



Note

The transmit interval value “3” is 3.3 msec. Also, the values 60000 and 600000 are supported in hardware.

Example

The following command changes the interval the UP MEP (previously configured on port 2:4) uses to send CCM messages on the 350 association in the finance domain to 10 seconds:

```
configure cfm domain finance association 350 ports 2:4 end-point up transmit-interval 10000
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure cfm domain association ports end-point ccm

```
configure cfm domain domain_name association association_name ports
port_list end-point [up | down ] ccm [disable | enable]
```

Description

This command is used to enable or disable sending CCMs on a given MEP.

Syntax Description

<i>domain_name</i>	Enter the domain associated with the MA you are configuring.
<i>association_name</i>	IEEE 802.1ag or ITU-T Y.1731 association name.
<i>port_list</i>	Enter the port number you want to configure as either an UP or DOWN MEP.

Default

Enabled.

Usage Guidelines

Each MA needs at least two MEPs that can reach each other to exchange CCM messages.



Note

Ensure that you assigned the port number correctly to the UP MEP and to the DOWN MEP, or the CCM messages go in the wrong direction.

These ports must already be in the MA (VLAN or VMAN) prior to assigning a MEP function to them. If you try to assign a port not in the MA as an end-point, the system returns the following message:

The following port(s) <portlist> are not part of the associations VLAN.

Use the `show cfm` command to verify your configuration.

Example

```
configure cfm domain "MD1" association "MD1v1" ports 17 end-point down delete group
"eapsCfmGrp"
```

History

This command was first available in ExtremeXOS 12.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure cfm domain association ports end-point mepid

```
configure cfm domain domain-name association association_name ports
port_list end-point [up | down] mepid mepid
```

Description

Allows you to change the MEP ID for a previously configured MEP. Each MEP within a single MA must have a unique MEP ID.

Syntax Description

<i>domain_name</i>	Enter the domain associated with the MA you are configuring.
<i>association_name</i>	IEEE 802.1ag or ITU-T Y.1731 association name.
<i>port_list</i>	Enter the port number you want to change the MEP ID.
up	Enter this variable if you are changing the MEP ID on an UP MEP.
down	Enter this variable if you are changing the MEP ID on a DOWN MEP.
<i>mepid</i>	Enter the new value for this MEP. The range is 1 to 8191. NOTE: On each MA, each MEPID must be unique.

Default

N/A.

Usage Guidelines

Use this command to change the MEPID on a previously configured UP or DOWN MEP. If you attempt to change the MEPID on a port that is either not an MEP or having wrong MEP type, the system returns an error message.

Use the `show cfm` command to verify your configuration.

Example

The following command changes the MEP ID to 75 on the previously configured port 2:4 UP MEP on the 350 association in the finance domain:

```
configure cfm domain finance association 350 ports 2:4 end-point up mepid 75
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure cfm domain association ports end-point sender-id-ipaddress

```
configure cfm domain domain_name association association_name ports
port_list end-point [up | down ] sender-id-ipaddress [disable |
enable ip-address]
```

Description

This command is used to disable or enable configuring the sender-id-ipaddress on a given MEP.

Syntax Description

<i>domain_name</i>	Enter the domain associated with the MA you are configuring.
<i>association_name</i>	IEEE 802.1ag or ITU-T Y.1731 association name.
<i>port_list</i>	Enter the port number.
<i>ip-address</i>	Specifies the IP address that is sent in the sender-id TLV of the CFM PDUs.

Default

Disable.

Usage Guidelines

Each MA needs at least two MEPs that can reach each other to exchange CCM messages.



Note

Ensure that you assigned the port number correctly to the UP MEP and to the DOWN MEP, or the CCM messages go in the wrong direction.

You must create the MEP for which the configuration is being made before changing the configuration. Otherwise, the following error message is displayed:

The following port(s) <portlist> are not part of the associations VLAN.

Use the `show cfm` command to verify your configuration.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure cfm domain association ports end-point

```
configure cfm domain domain_name association association_name ports
port_list end-point [up | down] [enable | disable]
```

Description

Enables or disables an MEP.

Syntax Description

<i>domain_name</i>	Specifies the domain name.
<i>association_name</i>	IEEE 802.1ag or ITU-T Y.1731 association name.
<i>port_list</i>	Specifies the ports to configure.
up	Specifies that the end point is up.
down	Specifies that the end point is down.

Default

MEP is enabled by default.

Usage Guidelines

Use this command to enable or disable an MEP.

Use the [show cfm](#) command to verify your configuration.

History

This command was first available in ExtremeXOS 12.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure cfm domain association remote-mep mac-address

```
configure cfm domain domain-name association association_name remote-mep  
    mepid mac-address mac_address
```

Description

Allows you to modify the MAC address of an existing MEP.

Syntax Description

<i>domain_name</i>	Enter the domain associated with the MA you are configuring.
<i>association_name</i>	IEEE 802.1ag or ITU-T Y.1731 association name.
<i>mepid</i>	Specifies the MEP ID of the remote MEP being modified. The range is 1 to 8191.
<i>mac_address</i>	Specifies the MAC address for the remote MEP being modified.

Default

N/A.

Usage Guidelines

Use this command to modify a remote MEP with given MEP ID and MAC address in an existing association. Use the [show cfm detail](#) command to verify your configuration.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure cfm domain delete association

```
configure cfm domain domain_name delete association association_name
```

Description

Deletes a maintenance association (MA), including all its configured values, from the switch.

Syntax Description

<i>domain_name</i>	Enter the domain associated with the MA you are deleting.
<i>association_name</i>	IEEE 802.1ag or ITU-T Y.1731 association name.

Default

N/A.

Usage Guidelines

When you delete an association, or MA, you also remove all its configured values from the switch. These values include all configured MEPs, MIPs, and static remote MEPs.

Example

The following command deletes the MA test, in the domain of brazil, from the switch, along with all its configured MIPs, MEPs, and static remote MEPs:

```
configure cfm domain brazil delete association test
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure cfm domain md-level

```
configure cfm domain domain_name md-level level
```

Description

Changes a previously configured MD level for the specified domain.

Syntax Description

<i>domain_name</i>	Enter the name of the domain for which you want to change the MD level.
<i>level</i>	Specifies the new MD level you are assigning to this domain. Enter a value between 0 and 7.

Default

N/A.

Usage Guidelines

You can have up to 8 domains on a switch, and each one must have a unique MD level. Thus, a given MD level exists only once on a switch.

The IEEE standard 801.2ag specifies different levels for different network users, as follows:

- 5 to 7 for end users
- 3 and 4 for Internet service providers (ISPs)
- 0 to 2 for operators (entities carrying the information for the ISPs)

Example

The following command changes the MD level of a previously created domain extreme to 2:

```
configure cfm domain extreme md-level 2
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure cfm group add rmep

```
configure cfm group group_name add rmep mepid
```

Description

This command allows you to create and associate an RMEP to a group.

Syntax Description

<i>mepid</i>	Specifies the MEP ID of the remote MEP being created. The range is 1 to 8191.
--------------	---

Default

N/A.

Usage Guidelines

Use this command to create and associate an RMEP to a group.

Example

```
configure cfm group "eapsCfmGroup" add rmep 2
```

History

This command was first available in ExtremeXOS 15.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure cfm group delete rmep

```
configure cfm group group_name delete rmep [mepid | all]
```

Description

This command allows you to delete one or all RMEPs from a group.

Syntax Description

<i>mepid</i>	Specifies the MEP ID of the remote MEP being created. The range is 1 to 8191.
--------------	---

Default

N/A.

Usage Guidelines

Use this command to delete one or all RMEPs from a group.

Example

```
configure cfm group "eapsCfmGroup" delete rmep 2
```

History

This command was first available in ExtremeXOS 15.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure cfm segment add domain association

```
configure cfm segment segment_name add domain domain_name association
association_name
```

Description

Adds a CFM domain and association to a CFM segment.

Syntax Description

<i>segment_name</i>	An alpha numeric string identifying the segment name.
<i>domain_name</i>	Specifies the IEEE 802.1ag maintenance domain.
<i>association_name</i>	IEEE 802.1ag or ITU-T Y.1731 association name.

Default

N/A.

Usage Guidelines

Use this command to add a CFM domain and an association to a CFM segment. It is used to enable DMM/DMR in the association that is configured in the CFM domain.

Example

The following command adds the domain cfm3 and the association as3 to the segment s2.

```
configure cfm segment s2 add domain cfm3 association as3
```

To delete the domain and/or association, use the command, [configure cfm segment delete domain association](#).

History

This command was first available in ExtremeXOS 12.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure cfm segment delete domain association

```
configure cfm segment segment_name delete domain association
```

Description

Deletes a CFM domain from a CFM segment.

Syntax Description

<i>segment_name</i>	An alpha numeric string identifying the segment name.
---------------------	---

Default

N/A.

Usage Guidelines

Use this command to delete a CFM domain from a CFM segment.

Example

The following command deletes the domain and association from the segment s2.

```
configure cfm segment s2 delete domain association
```

History

This command was first available in ExtremeXOS 12.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure cfm segment dot1p

```
configure cfm segment segment_name dot1p dot1p_priority
```

Description

Configures the priority for the segment.

Syntax Description

<i>segment-name</i>	An alpha numeric string identifying the segment name.
<i>dot1p_priority</i>	Priority value that is set in the DMM/DMR. The range is 0 to 7.

Default

The default is 6.

Usage Guidelines

Use this command to configure the dot1p priority that a DMM/DMR frame can get.

Example

The following example configures a dot1p priority of 3 for segment s2.

```
configure cfm segment s2 dot1p 3
```

History

This command was first available in ExtremeXOS 12.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure cfm segment frame-delay dot1p

```
configure cfm segment segment_name frame-delay dot1p dot1p_priority
```

Description

This command configures the class of service for a particular cfm segment. This value is used to fill the dot1p priority bit in the Ethernet header during transmission.

If the optional keyword **frame-delay** is not specified, the same value of Dot1p will be used for both DMM and LMM. The optional keyword allows configuring different values for DMM and LMM.

Syntax Description

<i>segment_name</i>	An alpha numeric string identifying the segment name.
<i>dot1p_priority</i>	Priority value that is set in the DMM/DMR. The range is 0 to 7.

Default

N/A.

Usage Guidelines

Use this command to configure the class of service for a particular cfm segment.

Example

```
configure cfm segment frame-delay dot1p 4
```

History

This command was first available in ExtremeXOS 12.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure cfm segment frame-delay window

```
configure cfm segment segment_name frame-delay window window_size
```

Description

This command is used to configure the window size for calculating the alarm/clear threshold values for DMM and Severely Errored Second (SES) threshold for LMM. This window size denotes the total number of recent frames for which the threshold values will be measured.

If the optional keyword **frame-delay** or **frame-loss** is not specified, the same value of window size will be used for both DMM and LMM. The optional keyword allows configuring values for DMM and LMM.

Syntax Description

<i>segment_name</i>	Alphanumeric string identifying the segment name.
frame-delay	Y.1731 Ethernet frame delay measurement.
<i>window_size</i>	Window size for delay measurement; number of frames 1-1800 to be used.

Default

60.

Usage Guidelines

Use this command to configure the window size for calculating the alarm/clear threshold values for DMM and Severely Errored Second (SES) threshold for LMM.

Example

```
configure cfm segment cs2 frame-delay window 1000
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure cfm segment frame-delay/frame-loss transmit interval

```
configure cfm segment segment_name {frame-delay | frame-loss} transmit-  
interval interval
```

Description

Configures the delay between two consecutive DMM/LMM frames.

Syntax Description

<i>segment_name</i>	Alphanumeric string identifying the segment name.
frame-delay	Y.1731 Ethernet frame delay measurement.
frame-loss	Y.1731 Ethernet frame loss measurement.
<i>interval</i>	Transmit interval in seconds, with a range of 1 to 90.

Default

N/A.

Usage Guidelines

Configures the delay between two consecutive DMM/LMM frames. The configured delay would be for both continuous and on-demand transmission. This command is optional, and if not configured, the default interval would be 10 seconds.

If the optional keyword frame-delay or frame-loss is not specified, the same value of transmit-interval will be used for both DMM and LMM. The optional keyword allows configuring different values for DMM and LMM.

Example

```
configure cfm segment cs2 frame-delay transmit-interval 10  
configure cfm segment cs2 frame-loss transmit-interval 10
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure cfm segment frame-loss consecutive

```
configure cfm segment segment_name frame-loss consecutive frames
```

Description

This command is used to configure the number of consecutive measurements to be used to determine the availability status of a CFM segment.

Syntax Description

<i>segment_name</i>	Alphanumeric string identifying the segment name.
frame-loss	Y.1731 Ethernet frame loss measurement.

Default

10.

Usage Guidelines

This configuration is optional.

Example

```
configure cfm segment cs2 frame-loss consecutive 10
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure cfm segment frame-loss dot1p

```
configure cfm segment segment_name frame-loss dot1p dot1p_priority
```

Description

This command configures the class of service for a particular cfm segment. This value is used to fill the dot1p priority bit in the Ethernet header during transmission.

If the optional keyword `frame-loss` is not specified, the same value of Dot1p will be used for both DMM and LMM. The optional keyword allows configuring different values for DMM and LMM.

Syntax Description

<i>segment_name</i>	An alpha numeric string identifying the segment name.
<i>dot1p_priority</i>	Priority value that is set in the DMM/DMR. The range is 0 to 7.

Default

N/A.

Usage Guidelines

Use this command to configure the class of service for a particular cfm segment.

Example

```
configure cfm segment frame-loss dot1p 4
```

History

This command was first available in ExtremeXOS 12.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure cfm segment frame-loss mep

```
configure cfm segment segment_name frame-loss [add|delete] mep mep_id
```

Description

This command is used to add/delete the local MEP for a given CFM segment.

Syntax Description

<i>segment_name</i>	Alphanumeric string identifying the segment name.
frame-loss	Y.1731 Ethernet frame loss measurement.

Default

N/A.

Usage Guidelines

The MEP with the given MEP ID should already be created in the system. The domain and association for the segment should be configured before executing this command. If the domain and association are not configured, the command throws an error.

Configuring of local MEP is mandatory to start the Frame Loss measurements.

Example

```
configure cfm segment cs2 add mep 3
configure cfm segment cs2 delete mep 3
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure cfm segment frame-loss ses-threshold

```
configure cfm segment segment_name frame-loss ses-threshold percent
```

Description

This command is used to configure the percentage of frames lost in a measurement period for it to be marked as SES (Severely Errored Second).

Syntax Description

<i>segment_name</i>	Alphanumeric string identifying the segment name.
ses	Severely errored second.
frame-loss	Y.1731 Ethernet frame loss measurement.

Default

30%.

Usage Guidelines

This configuration is optional.

Example

```
configure cfm segment cs2 frame-loss ses-threshold .02
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure cfm segment frame-loss window

```
configure cfm segment segment_name frame-loss window window_size
```

Description

This command is used to configure the window size for calculating the alarm/clear threshold values for DMM and Severely Errored Second (SES) threshold for LMM. This window size denotes the total number of recent frames for which the threshold values will be measured.

If the optional keyword frame-delay or frame-loss is not specified, the same value of window size will be used for both DMM and LMM. The optional keyword allows configuring values for DMM and LMM.

Syntax Description

<i>segment_name</i>	Alphanumeric string identifying the segment name.
frame-loss	Y.1731 Ethernet frame loss measurement.
<i>window_size</i>	Window size for loss measurement; number of frames 1-1800 to be used.

Default

1200.

Usage Guidelines

Use this command to configure the window size for calculating the alarm/clear threshold values for DMM and Severely Errored Second (SES) threshold for LMM.

Example

```
configure cfm segment cs2 frame-loss window 900
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure cfm segment threshold

```
configure cfm segment segment_name [alarm-threshold | clear-threshold]
                               value
```

Description

Configures the alarm threshold and clear threshold.

Syntax Description

alarm-threshold	Specifies the minimum threshold percentage.
clear-threshold	Specifies the maximum threshold percentage.
<i>value</i>	Specified the threshold percentage in a range of 1-99%.

Default

Alarm threshold is 10% of the total frames received during the current window.

Clear-threshold is 95% of the total frames received during the current window.

Usage Guidelines

Use this command to configure the alarm and clear threshold value for a CFM segment. Upon reaching the alarm threshold, an error message is generated and displayed once, and the state is maintained until the threshold reaches the clear threshold value.

This command is optional, and if not configured the default intervals are used.

Example

The following commands configure an alarm threshold of 15% and a clear-threshold of 90% for segment-first.

```
configure cfm segment segment-first alarm-threshold 15
configure cfm segment segment-first clear-threshold 90
```

History

This command was first available in ExtremeXOS 12.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure cfm segment timeout

```
configure cfm segment segment_name timeout msec
```

Description

Configures the timeout for a segment.

Syntax Description

<i>segment_name</i>	An alpha numeric string identifying the segment name.
<i>msec</i>	Specifies the number of milliseconds. The range is 1 to 65535.

Default

50 milliseconds.

Usage Guidelines

Use this command to configure the timeout value for the reception of a DMR frame. If a DMR frame is not received within this specified time, that frame is considered as an errored frame, and if the number of errored frames reaches the alarm threshold of the current window size, an alarm is generated.

This command is optional, and if not configured, timeout is set to the default.

Example

The following command configures a timeout value of 45 milliseconds for the s4 segment:

```
configure cfm segment s4 timeout 45
```

History

This command was first available in ExtremeXOS 12.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure cfm segment transmit-interval

```
configure cfm segment segment_name {frame-delay | frame loss}transmit-  
interval interval
```

Description

Configures the transmission interval of DMM frames.

Syntax Description

<i>segment_name</i>	An alpha numeric string identifying the segment name.
frame-delay	Y.1731 Ethernet Frame Delay Measurement.
frame loss	Y.1731 Ethernet Frame Loss Measurement.
<i>interval</i>	Specifies the transmit interval in seconds. The range is 1 to 90.

Default

10 seconds.

Usage Guidelines

Use this command to configure the delay between two consecutive DMM frames. The configured delay is for both continuous and on-demand transmission. This command is optional, and if not configured the default interval is used.

Example

The following example configures a transmission interval of 5 seconds for segment s2.

```
configure cfm segment s2 transmit-interval 5
```

History

This command was first available in ExtremeXOS 12.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure cfm segment window

```
configure cfm segment segment_name window size
```

Description

Configures the measurement window size.

Syntax Description

<i>segment_name</i>	An alpha numeric string identifying the segment name.
<i>size</i>	Specifies the number of frames to be used for delay measurement. The range is 1 to 1800.

Default

60 frames.

Usage Guidelines

Use this command to configure the window size to be used for calculating the threshold values. This window size denotes the total number of recent frames for which the threshold values are to be measured.

This is an optional command and if not configured, the lower of either the default value or the total number of frames sent is used.



Note

MEPs with intervals 3 and 10 cannot be created in this domain as the domain name format is of dns type.

Example

The following command configures the measurement window size for the CFM segment segment-first at 55:

```
configure cfm segment segment-first window 55
```

History

This command was first available in ExtremeXOS 12.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure cli

```
configure cli [{ lines height } {columns width }]
```

Description

This command configures the number of lines and columns for the current login session only.

Syntax Description

lines	Number of lines on the screen.
<i>height</i>	Height of the screen.
columns	Number of columns on the screen.
<i>width</i>	Width of the screen.

Default

N/A.

Usage Guidelines

The screen size specified takes effect over whatever screen size the session may have started with or whatever the current settings may be. If the terminal emulation supports dynamic resizing of the window, this will cause the size set by this command to be overridden. The command accepts either lines or columns or both in either order.

Example

The show management command has been enhanced to display the current screen size:

```
# show management
CLI idle timeout           : Enabled (20 minutes)
CLI max number of login attempts : 3
CLI max number of sessions  : 8
CLI paging                 : Enabled (this session only)
CLI space-completion       : Disabled (this session only)
CLI configuration logging   : Disabled
CLI password prompting only : Disabled
CLI scripting               : Disabled (this session only)
CLI scripting error mode    : Ignore-Error (this session only)
CLI persistent mode        : Persistent (this session only)
CLI prompting               : Enabled (this session only)
CLI screen/window size     : 80 Lines 256 Columns (this session only)
CLI refresh                 : Enabled
Telnet access               : Enabled (tcp port 23 vr all)
                           : Access Profile : not set
SSH Access                  : ssh module not loaded.
Web access                  : Enabled (tcp port 80)
                           : Access Profile : not set
Total Read Only Communities : 1
Total Read Write Communities : 1
RMON                        : Disabled
SNMP access                 : Enabled
                           : Access Profile : not set
SNMP Notifications          : Enabled
SNMP Notification Receivers : None
SNMP stats:      InPkts 0      OutPkts 0      Errors 0      AuthErrors
0
                  Gets 0      GetNexts 0      Sets 0      Drops 0
SNMP traps:      Sent 0      AuthTraps Enabled
SNMP inform:     Sent 0      Retries 0      Failed 0
```

History

This command was first available in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure cli journal

```
configure cli journal size size
```

Description

This command configures the size of the historical list (journal) of the most recently executed CLI commands.

Syntax Description

journal	List of the most recently executed CLI commands.
size <i>size</i>	Configures the size (number) of remembered commands. Range is 50 to 200 (default = 100).

Default

One hundred commands are preserved in the journal by default.

Usage Guidelines

The journal retains as many as 200 of the most recently executed commands along with the timestamp and user name. Commands are saved even after logging off, rebooting, or switch crashes.

To view the journal, use the `show cli journal` command.

Example

The following example sets the journal size to 150:

```
configure cli journal size 150
```

History

This command was first available in ExtremeXOS 22.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure cli max-failed-logins

```
configure cli max-failed-logins num-of-logins
```

Description

Establishes the maximum number of failed logins permitted before the session is terminated.

Syntax Description

<i>num-of-logins</i>	Specifies the maximum number of failed logins permitted; the range is 1 to 10.
----------------------	--

Default

The default is three logins.

Usage Guidelines

The value must be greater than 0; the range is 1 to 10.

Example

The following command sets the maximum number of failed logins to five:

```
configure cli max-failed-logins 5
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure cli max-sessions

```
configure cli max-sessions num-of-sessions
```

Description

Limits number of simultaneous CLI sessions on the switch.

Syntax Description

<i>num-of-sessions</i>	Specifies the maximum number of concurrent sessions permitted. The range is 1 to 16.
------------------------	--

Default

The default is eight sessions.

Usage Guidelines

The value must be greater than 0; the range is 1 to 16.

Example

The following command limits the number of simultaneous CLI sessions to ten:

```
configure cli max-sessions 10
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure cli mode

```
configure cli mode [persistent | non-persistent]
```

Description

Configures the persistent nature of command execution for non-persistent commands.

Syntax Description

persistent	Configures command execution to be persistent.
non-persistent	Configures command execution to be not persistent.

Default

The default mode is non-persistent.

Usage Guidelines

All ExtremeXOS commands can operate in persistent mode, and a subset of the ExtremeXOS command set can operate in non-persistent mode. Commands that are executed in persistent mode become part of the saved switch configuration that persists when the switch is rebooted. Commands that are executed in non-persistent mode configure temporary changes that are not saved in the switch configuration and do not persist when the switch is rebooted.

Most commands operate only in persistent mode. The subset of commands that operate in non-persistent mode are called non-persistent-capable commands. The Universal Port feature uses the non-persistent-capable commands to configure temporary changes that could create security issues if the switch were rebooted or reset. The use of non-persistent-capable commands in scripts and Universal Port profiles allows you to make temporary configuration changes without affecting the default configuration the next time the switch is started.

The `configure cli mode` command affects only the non-persistent-capable commands, which are listed in the Universal Port chapter in the [Switch Engine 32.3 User Guide](#). By default, all commands operate in persistent mode with the following exceptions:

- In Universal Port dynamic profiles, the non-persistent-capable commands operate in non-persistent mode unless preceded by the `configure cli mode persistent` command in the profile.
- In the CLI, CLI scripts, and static profiles, the non-persistent-capable commands operate in non-persistent mode only when preceded by the `configure cli mode non-persistent` command.

You can use the `configure cli mode persistent` command and the `configure cli mode non-persistent` command to change the mode of operation for non-persistent-capable commands multiple times within a script, profile, or configuration session.

Example

The following example sets command execution to be persistent:

```
configure cli mode persistent
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure cli mode scripting

```
configure cli mode scripting [abort-on-error | ignore-error]
```

Description

Configures the error handling process for CLI scripting on the switch.

Syntax Description

abort-on-error	Configures Cli scripts to be aborted if a CLI error occurs.
ignore-error	Configures the script to be executed when CLI errors occur.

Default

CLI: ignore-error Static profiles: abort-on-error Dynamic profiles: abort-on-error

Usage Guidelines

You can change the error-handling options within the scripts.

Example

The following command configures the switch to ignore syntax errors in CLI scripts:

```
configure cli mode scripting ignore-error
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure cli moved-keywords

```
configure cli moved-keywords [hide | show {no-help}]
```

Description

Controls how old keywords that have been moved and redefined appear in the CLI.

Syntax Description

cli	Configures aspects of the CLI.
moved-keywords	Selects CLI keywords that were moved or processing options that were renamed.
hide	Deprecates old-moved keywords to hide them from help display.
show	Shows old-moved keywords and corresponding redirection help text. (Default)
no-help	Shows old-moved keywords, but does not show redirection help text.

Default

By default, the **show** option is in effect.

Usage Guidelines

ExtremeXOS has evolved and incorporated many new features over time. During this development, CLI keywords have been introduced that are not logically organized or do not conform to the CLI format standards. This command provides a way to manage how old keywords that have been moved and redefined appear in the CLI.

The option you select with this command, and if you elect to hide commands, which version of ExtremeXOS version was running when the hide command was issued, appear in the output of the `show management`.

Example

The following example shows old commands and displays help text:

```
# configure cli moved-keywords show
```

History

This command was first available in ExtremeXOS 30.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure cli password prompting-only

```
configure cli password prompting-only [ on | off ]
```

Description

This command allows you to configure prompting (with no echo) for all passwords, secrets, or keys.

Syntax Description

prompting-only	Prompting is required when entering passwords, keys, and secrets. The default is off.
on	Enable the option.
off	Disable the option.

Default

Off.

Usage Guidelines

Use this command to configure prompting (with no echo) for all passwords, secrets, or keys. Each CLI command with password arguments will be modified to use the new mode (designated with

flags="prompting-only" in the CLI syntax attribute specification). Prompting must be handled in the action script for that command.

History

This command was first available in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure cli script path

```
configure cli script path path
```

Description

Creates a search path for the **run/load script filename** command.

.

Syntax Description

<i>path</i>	Defines the colon-separated list of directories to search for CLI scripts. Default is <code>./usr/local/cfg</code> .
-------------	--

Default

`./usr/local/cfg`

Usage Guidelines

This setting only applies to the current session. This command must be added to `exshrc.xsf` in order to be persistent.

Example

The following example configures the default script search path:

```
configure cli script path
```

History

This command was first available in ExtremeXOS 31.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure cli script timeout

```
configure cli script timeout timeout
```

Description

Configures the maximum time a script can run.

Syntax Description

<i>timeout</i>	Defines the timeout period in seconds.
----------------	--

Default

Regular script: no time limit default.xsf: 500 seconds autoexec.xsf: 500 seconds

Usage Guidelines

This command configures the maximum run time for all scripts, including default.xsf and autoexec.xsf, which are described in *Software Upgrade and Boot Options* section in the [Switch Engine 32.3 User Guide](#). If no timeout period is configured, regular scripts do not timeout, and the default.xsf and autoexec.xsf scripts time out after 500 seconds.

If a script does not finish running in the configured time, command execution stops and an error message is logged. If the timer expires while a command is executing, the command execution continues and all following commands are not executed.

If the timer command is executed inside a script, the timer is reset. If the command is issued more than once inside a script the last timer command executed resets the timer. The timer is valid only for that session. The use of nested scripts does not extend the execution period. When the parent script reaches the timeout value, the parent script and all nested scripts terminate.

To configure a different timeout value for autoexec.xsf or default.xsf, the configure cli script timeout command should be the first command in the script.

When a script timeout value is configured, the following variables are created: \$CLI.SCRIPT_TIMEOUT and \$CLI.SCRIPT_TIME_REMAINING. If no timeout value is configured for a session, the variables are not created.

You can use the \$CLI.SCRIPT_TIMEOUT variable to adjust the timeout value. The \$CLI.SCRIPT_TIME_REMAINING variable returns the time remaining. When a timeout value is configured, the variable values are as follows:

- If no script is running, both \$CLI.SCRIPT_TIME_REMAINING and \$CLI.SCRIPT_TIMEOUT show the configured timeout value.
- If a script is aborted due to timeout, the \$CLI.SCRIPT_TIME_REMANING variable returns the value0.
- If a script finishes execution (before the timeout value is reached) the \$CLI.SCRIPT_TIME_REMANING variable returns the remaining time.

Example

The following example configures the switch to terminate a script after 120 seconds:

```
configure cli script timeout 120
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure cos-index

```
configure cos-index cos_index [{ qosprofile qosprofile } {ingress-meter
ing_meter } {replace-tos tos_value {mask tos_mask}}]
```

Description

This command is used to configure the *CoS (Class of Service)* index, which is used to assign *QoS (Quality of Service)* rate-shaping, rate-limiting, flood control, and 802.1p.

Syntax Description

<i>cos_index</i>	Class of Service (CoS) index value, range 0 - 255.
qosprofile	QoS profile.
<i>qosprofile</i>	QoS profile name.
ingress-meter	Ingress rate-limiter meter.
<i>ing_meter</i>	Ingress rate-limiter meter name.
replace-tos	Replace TS value.
<i>tos_value</i>	TOS replacement value.
mask	TOS replacement mask.
<i>tos_mask</i>	TOS replacement mask value.

Default

N/A.

Usage Guidelines

The CoS index (0-255) is used to assign QoS rate-shaping, rate-limiting, flood control, and 802.1p. The TOS value can be a value from 0-255. The TOS mask option allows for only certain bits of the field, those masked, to be change. If the mask is not specified in the ToS input, all bits are overwritten. The replace-dot1p value cannot be set for CoS indexes 0-7.

For indexes 0-7, the **replace-tos** option for the cos-index command will map to the `configure diffserv` commands, which are associated with the qosprofile, assigned through the `configure dot1p` command. Note that diffserv only replaces bits 0-5 of the TOS byte. Therefore, the replace-tos mask is fixed to 0xfc for cos-index 0-7 and the equivalent diffserv replace value is shifted left 2 bits. On some platforms, the hardware only allows replacement of bits 0-5. In which case, the mask is fixed to 0xfc and will result in an error if the user tries to change the mask.

Example

```
configure cos-index 51 qosprofile qp2 ingress-meter ingmeter2 replace-tos 64
```

History

This command was first available in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure database add server

```
configure database database_name add server [host_name | ip_address]
    {port port_number} {password [encrypted encrypted_password |
    password ]}
```

Description

Adds a server to an Automation Edge remote VXLAN Network Identifier (VNI)-device database.

Syntax Description

database	Adds a server to a remote VNI-database.
<i>database_name</i>	Adds a server to the named database.
add	Adds a server to the database.
server	Adds a server to the database.
<i>host_name</i>	Provides hostname of the remote database server.
<i>ip_address</i>	Provides the IP address of the remote server.
port	Configures the server TCP port number.
<i>port_number</i>	Specifies the server TCP port number. The default is 6,379.
password	Specifies providing a password for the database.
encrypted	Specifies providing an encrypted password for the database.
<i>encrypted_password</i>	Provides the encrypted password.
<i>password</i>	Provides the password.

Default

If not specified, the server TCP port number is 6,379.

Usage Guidelines

N/A.

Example

The following example adds a server at location 1.1.1.2 with password "secretpassword" to the database "database1":

```
# configure database database1 add server 1.1.1.2 password secretpassword
```

History

This command was first available in ExtremeXOS 31.1 as a demonstration feature.

This command is fully supported in ExtremeXOS 31.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure database delete server

```
configure database database_name delete server [host_name | ip_address | all]
```

Description

Deletes a server from an Automation Edge remote VXLAN network identifier (VNI)-device database.

Syntax Description

database	Deletes server from a remote VNI-database.
<i>database_name</i>	Deletes server from the named database.
delete	Deletes a server from the database.
server	Deletes a from to the database.
<i>host_name</i>	Specifies the hostname of the server to delete from the database.
<i>ip_address</i>	Specifies the IP address of the server to delete from the database.
all	Specifes deleting all servers from the database.

Default

N/A.

Usage Guidelines

N/A.

Example

The following example deletes a server at location 1.1.1.2 from database "database1":

```
# configure database database1 delete server 1.1.1.2
```

History

This command was first available in ExtremeXOS 31.1 as a demonstration feature.

This command is fully supported in ExtremeXOS 31.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure database max-retry-interval

```
configure database max-retry-interval retry_interval
```

Description

Specifies the maximum value for exponentially increasing time interval between retries for an Automation Edge remote VXLAN network identifier (VNI)-device database.

Syntax Description

database	Makes time interval retries remote VNI-database.
max-retry-interval	Specifies setting the maximum value for exponentially increasing time interval between retries.
<i>retry_interval</i>	Specifies the value for the maximum time interval between retries in seconds. The default is 600. The range is 1 to 3,600.

Default

If not specified, the maximum retry interval is 600 seconds.

Usage Guidelines

N/A.

Example

The following example sets the maximum retry interval to 800 seconds:

```
# configure database max-retry-interval 800
```

History

This command was first available in ExtremeXOS 31.1 as a demonstration feature.

This command is fully supported in ExtremeXOS 31.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure database server password

```
configure database database_name server [host_name | ip_address]
      password [encrypted encrypted_password | password ]
```

Description

Updates an existing Automation Edge server in a remote VXLAN Network Identifier (VNI)-device database with a new password.

Syntax Description

database	Specifies the server of a database for a password change.
<i>database_name</i>	Specifies the server of the named database for a password change.
server	Specifies the server for a password change.
<i>host_name</i>	Provides hostname of the remote database server.
<i>ip_address</i>	Provides the IP address of the remote server.
password	Specifies providing a password for the database.
encrypted	Specifies providing an encrypted password for the database.
<i>encrypted_password</i>	Provides the encrypted password.
<i>password</i>	Provides the password.

Default

N/A.

Usage Guidelines

N/A.

Example

The following example changes the password for the server at location 1.1.1.2 to "differentpassword" in the database "database1":

```
# configure database database1 server 1.1.1.2 password differentpassword
```

History

This command was first available in ExtremeXOS 31.1 as a demonstration feature.

This command is fully supported in ExtremeXOS 31.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure debug core-dumps

```
configure debug core-dumps [ off | directory_path ]
```

Description

Enables or disables the sending of core dump files to the internal memory or a USB 2.0 storage device.

Syntax Description

off	Specifies that the switch does not save core dump files to memory or to removable storage devices.
<i>directory_path</i>	Directory path (USB is <code>/usr/local/ext</code> ; internal memory is <code>/usr/local/tmp</code> (default); and home directory is <code>/usr/local/cfg</code>).

Default

Beginning with ExtremeXOS 11.6, core dumps to internal memory (`/usr/local/tmp`) is enabled by default.

Usage Guidelines



Note

Use this command only under the guidance of Extreme Networks Technical Support personnel to troubleshoot the switch.

The switch only generates core dump files and writes them to the specified device in the following situations:

- If an ExtremeXOS process fails.
- When forced under the guidance of Extreme Networks Technical Support.

If you configure the switch to write core dump files to the internal memory and attempt to download a new software image, you might have insufficient space to complete the image download. If this occurs, move or delete the core dump files from the internal memory. For example, if the switch supports a removable storage device that has space available, transfer the files to the device. On switches without removable storage devices, transfer the files from the internal memory card to a TFTP server. This frees up space on the internal memory card while keeping the core dump files.

Before you can enable and save debug information to a removable storage device, you must install the device. For more information about installing a removable storage device, refer to the hardware documentation.

After you use the `eject usb-device` command and manually remove a USB device, you are prompted to select another location to write the debug files to.

Stackables in Stack Mode

This command works only from the master node. If you enable it on stack master, it is applicable for all nodes.

Example

The following example enables a switch to save debug information to a USB device:

```
configure debug core-dumps /usr/local/ext
```

The following example enables the switch to save debug information to internal memory:

```
configure debug core-dumps /usr/local/tmp
```

History

This command was first available in ExtremeXOS 11.1.

The internal-memory parameter was added in ExtremeXOS 11.2.

Support for USB 2.0 storage devices was added in ExtremeXOS 12.5.3.

The options **memorycard** and **internal-memory** were removed and the variable *directory_path* was added in ExtremeXOS 30.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure dhcp ipv6 client identifier-type

```
configure dhcp ipv6 client identifier-type [ link-layer {plus-time} |  
      vendor-specific
```

Description

This command configures the DHCPv6 client identifier type for the client. A *DHCP* server uses this identifier-type to identify clients for the selection of configuration parameters.

Syntax Description

<code>dhcp</code>	Configure DHCP
<code>ipv6</code>	Configure DHCP IPv6 client
<code>client</code>	Configure DHCP IPv6 client
<code>identifier-type</code>	Configure DHCP IPv6 client identifier type
<code>link-layer</code>	Configure link-layer address (system MAC) as DHCP IPv6 client identifier
<code>plus-time</code>	Configure link-layer address plus current time as DHCP IPv6 client identifier
<code>vendor-specific</code>	Configure DHCP IPv6 client identifier by prepending the vendor-specific IANA value

Default

IPv4.

Usage Guidelines

Use this command to configure the DHCPv6 client identifier type for the client.

History

This command was first available in ExtremeXOS 15.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure diagnostics privilege

```
configure diagnostics privilege [admin | user]
```

Description

This command configures the user privilege level needed to view diagnostic results.

Syntax Description

privilege	Configure minimum privilege level needed to view diagnostic results.
admin	Only admin (read-write) accounts can view diagnostic results.
user	User (read-only) accounts can view diagnostic results also (default).

Default

User.

Usage Guidelines

Use this command to configure the privilege level required to view diagnostic results.

History

This command was first available in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure diffserv examination code-point qosprofile

```
configure diffserv examination code-point code_point {qosprofile}  
                                qosprofile
```

Description

Configures the default ingress DiffServ code point (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). Supported values are 0 to 63.
qosprofile	Specifies the QoS profile to which the DiffServ code point is mapped.

Default

See [Table 4](#) below.

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 shown in the following table.

Table 4: Default DiffServ Code Point-to-QoS Profile Mapping

Code Point	ExtremeSwitching Series Switches QoSProfile
0-7	QP1
8-15	QP1
16-23	QP1
24-31	QP1
32-39	QP1
40-47	QP1
48-55	QP1
56-63	QP8

Example

The following command specifies that code point 25 be assigned to QP2:

```
# configure diffserv examination code-point 25 qosprofile qp2
```

History

This command was first available in ExtremeXOS 11.0.

The ports keyword was first available in ExtremeXOS 12.2.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure diffserv replacement code-point

```
configure diffserv replacement [{qosprofile} qosprofile | priority
priority] code-point code_point
```

Description

Configures the egress Diffserv replacement mapping for either a [QoS](#) profile or an 802.1p priority value.

Syntax Description

qosprofile	Specifies a QoS profile.
priority	Specifies an 802.1p priority value to map to a code point.
code_point	Specifies a 6-bit value to be used as the replacement DSCP in the IPv4 or IPv6 header.

Default

N/A.

Usage Guidelines



Note

We recommend that you use the qosprofile *qosprofile* value to configure this parameter.

Egress packets contain the DSCP assigned to the QoS profile, which can be selected by the 802.1p code point or by an [ACL](#). The default 802.1p priority value to QoS profile to DSCP mapping is shown in the following table.

Table 5: Default QoS Profile-to-802.1p Priority Value-to-Code Point

802.1p Priority Value	ExtremeSwitching Series Switches QoS Profile	DSCP
0	QP1	0
1	QP1	8
2	QP1	16
3	QP1	24
4	QP1	32
5	QP1	40
6	QP1	48
7	QP8	56

Example

The following command specifies that a code point value of 5 should be used to replace the DiffServ (TOS) bits in packets in QP2:

```
# configure diffserv replacement qosprofile qp2 code-point 5
```

History

This command was first available in ExtremeXOS 11.0.

The ports keyword was first available in ExtremeXOS 12.2.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure dns cache analytics [add | delete] protected-client

```
configure dns cache analytics [add | delete]protected-client [client_ip
    netmask | ipNetmask] [{vr} vr_name]
```

Description

Configures the protected client list for the Domain Name System (DNS) cache analytics for the virtual router (VR).

Syntax Description

dns	Domain Name System.
cache	Specifies configuring DNS cache.
analytics	Specifies configuring DNS cache analytics.
add	Specifies adding to the protected client list.
delete	Specifies deleting from the protected client list.
protected-client	Specifies configuring the protected client list.
<i>client_ip</i>	Specifies the IPv4 network address of the protected client.
<i>netmask</i>	Specifies the IP address netmask of the protected client.
<i>ipNetmask</i>	Specifies the IP address/mask length of the protected client.
vr	Specifies the VR.
<i>vr_name</i>	Specifies the VR name. If not specified, the VR of the current command context is used.

Default

If not specified, by default the VR of the current command context is used.

Usage Guidelines

Administrators can use this command to restrict the collection of DNS analytics for a protected client. When you configure the client IP subnet in the protected list, DNS queries from configured protected clients are erased from the analytics database and future queries are not stored.

Example

The following example adds the client at IP address 192.168.3.3 on VR-Default to the protected client list:

```
# configure dns cache analytics add protected-client 192.168.3.3 255.255.255.255 VR-Default
```

or

```
# configure dns cache analytics add protected-client 192.168.3.3/32 VR-Default
```

The following example adds the subnet 192.168.3.0 on VR-Default to the protected client list:

```
# configure dns cache analytics add protected-client 192.168.3.0 255.255.255.0 VR-Default
```

or

```
# configure dns cache analytics add protected-client 192.168.3.0/24 VR-Default
```

The following example removes the client 192.168.3.3 on VR-Default from the protected client list:

```
# configure dns cache analytics delete protected-client 192.168.3.3 255.255.255.255 VR-Default
```

or

```
# configure dns cache analytics delete protected-client 192.168.3.3/32 VR-Default
```

History

This command was first available in ExtremeXOS 30.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure dns cache add | delete name-server

```
configure dns cache [add | delete] name-server ip_address [{vr} vr_name]
```

Description

Adds or deletes a Domain Name System (DNS) name server.

Syntax Description

<i>dns</i>	Domain Name System.
cache	Specifies adding or deleting DNS name server.
add	Specifies adding a name server.
delete	Specifies deleting a name server.
name-server	Specifies adding or deleting a DNS name server.
<i>ip_address</i>	Specifies the IP address of the DNS name server.
vr	Configures the VR on which the DNS name server is accessible.
<i>vr_name</i>	Specifies the VR on which the DNS name server is accessible. If not specified, the VR of the current command context is used.

Default

If no VR name is specified, the VR of the current command context is used.

Usage Guidelines

You can configure a maximum of 8 name servers.

To view the current DNS name servers, use the command `show dns cache name-server`.

Example

The following example adds a DNS name server located at 1.1.1.2:

```
# configure dns cache add name-server 1.1.1.2
```

History

This command was first available in ExtremeXOS 30.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure dns cache analytics

```
configure dns cache analytics [{timeout minutes} {max-entries  
max_entries}]
```

Description

Configures Domain Name System (DNS) cache analytics.

Syntax Description

dns	Domain Name System.
cache	Specifies DNS cache.
analytics	Specifies configuring DNS analytics.
timeout	Specifies setting the timeout period for analyzed DNS queries. After this time, existing entries are flushed.
<i>minutes</i>	Specifies the timeout value in minutes. The range is 1 to 1,440. The default is 1,440.
max-entries	Specifies the maximum number of analyzed DNS queries in the database. When this limit is met, new entries start replacing old entries.
<i>max_entries</i>	Specifies the value for the maximum analyzed queries. The range is 1,000 to 10,000. The default is 10,000.

Default

The default for the timeout period is 1,440 minutes.

The default for the maximum number of entries is 10,000.

Usage Guidelines

If query Q1 is learned at time t1 and the timeout period is configured as 5 minutes, this entry is removed within t1 + 5 minutes.

To manually clear the DNS cache analytics, use the command `clear dns cache analytics entries {{vr} vr_name}`.

Example

The following example sets the maximum number of entries to 2,000:

```
# configure dns cache analytics max-entries 2000
```

The following example sets the timeout period to 500 minutes:

```
# configure dns cache analytics timeout 500
```

History

This command was first available in ExtremeXOS 30.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure dns-client add

```
configure dns-client add [domain-suffix domain_name | name-server
ip_address {vr vr_name}]
```

Syntax Description

domain-suffix	Specifies adding a domain suffix.
<i>domain_name</i>	Specifies a domain name.
name-server	Specifies adding a name server.
<i>ip_address</i>	Specifies an IP address for the name server.
vr	Specifies use of a virtual router. Note: User-created VRs are supported only on the platforms listed for this feature in the Switch Engine 32.3 Feature License Requirements document..
<i>vr_name</i>	Specifies a virtual router.

Description

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

Default

N/A.

Usage Guidelines

The domain suffix list can include up to six items.

If the use of all previous names fails to resolve a name, the most recently added entry on the domain suffix list will be the last name used during name resolution. 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.

Up to eight DNS name servers can be configured. The default value for the virtual router used by the DNS client option is *VR-Default*.

Example

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

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

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

```
configure dns-client add name-server 10.1.2.1
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure dns-client default-domain

```
configure 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

<i>domain_name</i>	Specifies a default domain name.
--------------------	----------------------------------

Default

N/A.

Usage Guidelines

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 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:

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

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure dns-client delete

```
configure dns-client delete [domain-suffix domain_name | name-server  
ip_address {vr vr_name}]
```

Description

Deletes a domain suffix from the domain suffix list or a name server from the available server list for the DNS client.

Syntax Description

domain-suffix	Specifies deleting a domain suffix.
<i>domain_name</i>	Specifies a domain name.
name-server	Specifies deleting a name server.
<i>ip_address</i>	Specifies an IP address for the name server.

vr	Specifies deleting a virtual router. Note: User-created VRs are supported only on the platforms listed for this feature in the Switch Engine 32.3 Feature License Requirements document.
<i>vr_name</i>	Specifies a virtual router.

Default

N/A.

Usage Guidelines

Specifying a domain suffix 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.

The default value for the virtual router used by the DNS client option is [VR-Default](#).

Example

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

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

The following example removes a DNS server from the list:

```
configure dns-client delete name-server 10.1.2.1
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure dos-protect acl-expire

```
configure dos-protect acl-expire seconds
```

Description

Configures the denial of service protection [ACL](#) expiration time.

Syntax Description

<i>seconds</i>	Specifies how long the ACL is in place.
----------------	---

Default

The default is 5 seconds.

Usage Guidelines

This command configures how long the DoS protection ACL remains in place.

Example

This example sets the ACL expiration time to 15 seconds:

```
configure dos-protect acl-expire 15
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure dos-protect interval

```
configure dos-protect interval seconds
```

Description

Configures the denial of service protection interval.

Syntax Description

<i>seconds</i>	Specifies how often the DoS protection counter is monitored.
----------------	--

Default

The default is one second.

Usage Guidelines

This command configures how often the DoS protection counter is monitored.

Example

This example sets the interval to 5 seconds:

```
configure dos-protect interval 5
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure dos-protect trusted ports

```
configure dos-protect trusted-ports [ports [ports | all] | add-ports
[ports-to-add | all] | delete-ports [ports-to-delete | all]]
```

Description

Configures the list of trusted ports.

Syntax Description

<i>ports</i>	Specifies the trusted ports list.
<i>ports-to-add</i>	Specifies the ports to add to the trusted ports list.
all	Specifies all the ports.
<i>ports-to-delete</i>	Specifies the ports to delete from the trusted ports list.

Default

N/A.

Usage Guidelines

Traffic from trusted ports will be ignored when DoS protect counts the packets to the CPU. If we know that a machine connected to a certain port on the switch is a safe "trusted" machine, and we know that we will not get a DoS attack from that machine, the port to which this machine is connected can be configured as a trusted port, even though a large amount of traffic is going through this port.

Example

This example sets the trusted port list to 3:1-3:7:

```
configure dos-protect trusted-ports ports 3:1-3:7
```

This example adds the trusted port 3:8 to the current list (use this command with a network administrator machine not connected to the internet that is attached to port 3:8):

```
configure dos-protect trusted-ports add-ports 3:8
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure dos-protect type l3-protect alert-threshold

```
configure dos-protect type l3-protect alert-threshold packets
```

Description

Configures the denial of service protection alert threshold.

Syntax Description

<i>packets</i>	Specifies how many packets in an interval will cause an alert.
----------------	--

Default

The default is 4000 packets.

Usage Guidelines

This command configures how many packets received in an interval will cause a DoS protection alert. When an alert occurs, the packets are analyzed, and a temporary ACL is applied to the switch.

Example

This example sets the alert threshold to 8000 packets:

```
configure dos-protect type l3-protect alert-threshold 8000
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure dos-protect type l3-protect notify-threshold

```
configure dos-protect type l3-protect notify-threshold packets
```

Description

Configures the denial of service protection notification threshold.

Syntax Description

<i>packets</i>	Specifies how many packets in an interval will cause a notification.
----------------	--

Default

The default is 3500 packets.

Usage Guidelines

This command configures how many packets received in an interval will cause a DoS protection notification.

Example

This example sets the notification threshold to 7500 packets:

```
configure dos-protect type l3-protect notify-threshold 7500
```

History

This command was first available in ExtremeXOS 11.1

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure dot1p type

```
configure dot1p type dot1p_priority {qosprofile} qosprofile {ingress-meter [ ing_meter | none ]}
```

Description

Configures an 802.1p priority to [QoS](#) profile mapping for the specified ports.

Syntax Description

<i>dot1p_priority</i>	Specifies the 802.1p priority value. The value is an integer between 0 and 7.
<i>qosprofile</i>	Specifies a specific QoS profile. The value range is QP1 to QP8.
ingress-meter	Ingress rate-limiter meter.

<i>ing_meter</i>	Ingress rate-limiter meter name.
none	Dot1p examination rule has no ingress-meter (default if ingress-meter is unspecified).

Default

The default mapping of each 802.1p priority value to QoS profile is shown in the following table.

Table 6: Default 802.1p Priority Value-to-QoS Profile Mapping

802.1p Priority Value	ExtremeSwitching Series Switches Default QoS Profile
0	QP1
1	QP1
2	QP1
3	QP1
4	QP1
5	QP1
6	QP1
7	QP8

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.

You must create the QoS profile first, using the `create qosprofile [QP2| QP3 | QP4 | QP5 | QP6 | QP7]` command, to map the 802.1p information to QoS profile 2 through 7.

SummitStack Only

You must create the QoS profile first, using the `create qosprofile [QP2| QP3 | QP4 | QP5 | QP6 | QP7]` command, to map the 802.1p information to QoS profile 2 through 6. You cannot create QP7 in a SummitStack.

Example

The following commands reassign (from the default) the QoS profiles associated with 802.1p priority values 1 and 2:

```
# configure dot1p type 2 qosprofile qp2
# configure dot1p type 1 qosprofile qp3
```

The following examples use the **ingress-meter** option:

```
# configure dot1p type 1 qosprofile qp5 ingress-meter ingmeter0
# configure dot1p type 2 qp3 ingress-meter ingmeter2
# configure dot1p type 3 qp4
```

History

This command was first available in ExtremeXOS 11.0.

The **ingress-meter**, *ing_meter*, and **none** options were added in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure eaps add control vlan

```
configure eaps name add control {vlan} vlan_name
```

Description

Adds the specified control VLAN to the specified EAPS domain.

Syntax Description

<i>name</i>	Specifies the name of an EAPS domain.
<i>vlan_name</i>	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 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 can be added as members of the control VLAN.
- The ring ports of the control VLAN must be tagged.

A control VLAN cannot belong to more than one EAPS domain. When the EAPS domain is active, you cannot delete or modify the configuration of the control VLAN.

By default, EAPS protocol data units (PDUs) are automatically assigned to QoS profile QP8. This ensures that the control VLAN messages reach their intended destinations. You do not need to configure a QoS profile for the control VLAN.

The VLAN must already exist before you can add it as a control VLAN. If you attempt to add a VLAN that does not exist, the switch displays a message similar to the following:

```
* Switch.8 # configure eaps megtest add control foo^%% Invalid input detected at '^'
marker.
```

To create the VLAN, use the `create vlan` command.

Example

The following command adds the control VLAN keys to the EAPS domain `eaps_1`.

```
configure eaps eaps_1 add control vlan keys
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure eaps add protected vlan

```
configure eaps name add protected {vlan} vlan_name
```

Description

Adds the specified protected VLAN to the specified EAPS domain.

Syntax Description

<i>name</i>	Specifies the name of an EAPS domain.
<i>vlan_name</i>	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.

A protected VLAN can be added to one or more EAPS domains.

When you configure 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.

The VLAN must already exist before you can add it as a protected VLAN. If you attempt to add a VLAN that does not exist, the switch displays a message similar to the following:

```
* Switch.5 # configure eaps megtest add protected foo^%% Invalid input detected at '^'
marker.
```


To create the VLAN, use the `create vlan` command.

Example

The following command adds the protected VLAN orchid to the EAPS domain eaps_1:

```
configure eaps eaps_1 add protected vlan orchid
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure eaps cfm

```
configure eaps cfm [add | delete] group group_name
```

Description

Notifies the CFM that EAPs is interested in notifications for the specified MEP and RMEP pair.

Syntax Description

cfm	Connectivity Fault Management.
add	Add a MEP group.
delete	Delete a MEP group.
group <i>group_name</i>	MEP group to bind.

Default

N/A.

Usage Guidelines

This command notifies CFM that EAPs is interested in notifications for this MEP and RMEP pair. This MEP should already be bound to a physical port, so when notification is received, EAPS associates that notification with a ring-port failure.

Example

The following command deletes the control VLAN keys from the EAPS domain eaps_1:

```
configure eaps cfm add
```

History

This command was first available in ExtremeXOS 15.2.

Platform Availability

This command is available on all ExtremeXOS platforms; however, not all platforms support hardware-based CFM. Platforms with no hardware-based CFM support are limited to software-based CFM transmit intervals of 100 ms or higher. Hardware-based intervals can go as low as 3.3 ms.

configure eaps config-warnings off

```
configure eaps config-warnings off
```

Description

Disables the loop protection warning messages displayed when configuring specific *EAPS* parameters.

Syntax Description

This command has no arguments or variables.

Default

By default, loop protection warnings are enabled and displayed when configuring specific EAPS parameters.

Usage Guidelines

This is a global EAPS command. You configure the warning message display on a per switch basis, not per EAPS domain.

When configuring the following EAPS parameters, the switch displays loop protection warning messages:

- Adding EAPS primary or secondary ring ports to a *VLAN*
- Deleting a protected VLAN
- Disabling the global EAPS setting on the switch
- Disabling an EAPS domain
- Configuring an EAPS domain as a transit node
- Unconfiguring EAPS primary or secondary ring ports from an EAPS domain

We recommend that you keep the loop protection warning messages enabled. If you have considerable knowledge and experience with EAPS, you might find the EAPS loop protection warning messages unnecessary. For example, if you use a script to configure your EAPS settings, disabling the warning messages allows you to configure EAPS without replying to each interactive yes/no question.

To confirm the setting on the switch, use the following command:

```
show eaps {eapsDomain} {detail}
```

Example

The following command disables the loop protection warning messages:

```
configure eaps config-warnings off
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure eaps config-warnings on

```
configure eaps config-warnings on
```

Description

Enables the loop protection warning messages displayed when configuring specific *EAPS* parameters.

Syntax Description

This command has no arguments or variables.

Default

By default, loop protection warnings are enabled and displayed when configuring specific EAPS parameters.

Usage Guidelines

This is a global EAPS command. You configure the warning message display on a per switch basis, not per EAPS domain.

When configuring the following EAPS parameters, the switch displays loop protection warning messages:

- Adding EAPS primary or secondary ring ports to a *VLAN*
- Deleting a protected VLAN
- Disabling the global EAPS setting on the switch
- Disabling an EAPS domain
- Configuring an EAPS domain as a transit node
- Unconfiguring EAPS primary or secondary ring ports from an EAPS domain

We recommend that you keep the loop protection warning messages enabled.

Example

The following command enables the loop protection warning messages:

```
configure eaps config-warnings on
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure eaps delete control vlan

```
configure eaps name delete control {vlan} vlan_name
```

Description

Deletes the specified control VLAN from the specified EAPS domain.

Syntax Description

<i>name</i>	Specifies the name of an EAPS domain.
<i>vlan_name</i>	Specifies the name of the control VLAN.

Default

N/A.

Usage Guidelines

None.

Example

The following example deletes the control VLAN keys from the EAPS domain eaps_1:

```
configure eapseaps_1 delete control vlan keys
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure eaps delete protected vlan

```
configure eaps name delete protected {vlan} vlan_name
```

Description

Deletes the specified protected *VLAN* from the specified *EAPS* domain.

Syntax Description

<i>name</i>	Specifies the name of an EAPS domain.
<i>vlan_name</i>	Specifies the name of the protected VLAN.

Default

N/A.

Usage Guidelines

To prevent loops in the network, you must delete the ring ports (the primary and the secondary ports) from the protected VLAN **before** deleting the protected VLAN from the EAPS domain. Failure to do so can cause a loop in the network.

The switch displays by default a warning message and prompts you to delete the VLAN from the EAPS domain. When prompted, do one of the following:

- Enter *y* delete the VLAN from the specified EAPS domain.
- Enter *n* or press **[Return]** to cancel this action.

If you have considerable knowledge and experience with EAPS, you might find the EAPS loop protection warning messages unnecessary. For more information, see the [configure eaps config-warnings off](#) command.

Useful show Commands

Use the following show commands to display information about your EAPS domain, including protected VLANs and primary and secondary ports:

- [show vlan](#)—This command displays summary information for all of the VLANs on the device. If the VLAN is a protected VLAN, the P flag appears in the flag column. To see more detailed information about the protected VLAN, use the following command: [show vlanvlan_name](#) .
- [show eaps](#)—This command displays summary EAPS domain information, including the name of the domain and the primary and secondary ports. To see more detailed information, including the name of the protected VLAN and the primary and secondary ports, use the [show eapseapsDomain](#) command.

- `show vlan eaps`—This command displays whether the VLAN is a control or partner VLAN for an EAPS domain. This command also displays if the VLAN is not a member of any EAPS domain.

Example

The following example deletes the protected VLAN orchid from the EAPS domain eaps_1:

```
configure eapseaps_1delete protected vlan orchid
```

The switch displays the following warning message and prompts you to confirm this action:

```
WARNING: Make sure EAPS ring-ports are deleted from the VLAN first. Otherwise deleting
the VLAN from the EAPS domain could cause a loop in the network! Are you sure you want to
remove the VLAN before deleting EAPS ring-ports.? (y/n)
```

Enter y to delete the VLAN from the specified EAPS domain. Enter n to cancel this action.

History

This command was first available in ExtremeXOS 11.0.

The interactive messages were added in ExtremeXOS 11.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure eaps failtime expiry-action

```
configure eaps name failtime expiry-action [open-secondary-port | send-
alert]
```

Description

Configures the action taken when the failtimer expires.

Syntax Description

<i>name</i>	Specifies the name of an <i>EAPS</i> domain.
open-secondary-port	Specifies to open the secondary port when the failtimer expires.
send-alert	Specifies that a critical message is sent to the syslog when the failtimer expires.

Default

Default is send-alert.

Usage Guidelines

By default the action is to send an alert if the failtimer expires. Instead of going into a Failed state, the master node remains in a Complete or Init state, maintains the secondary port blocking, and writes a critical error message to syslog warning the user that there is a fault in the ring. An [SNMP](#) trap is also sent.

If the EAPS ring contains non-EAPS devices, you must use the open-secondary-port parameter.



Note

Use caution when setting the failtimer expiry action to open-secondary port. Using this configuration, if the master node loses three consecutive hello PDUs, the failtimer expires—but there might not be a break in the ring. Opening the secondary port in this situation creates a loop.

Example

The following command configures the failtimer expiry action for EAPS domain eaps_1:

```
configure eapseaps_1 failtimeexpiry-action open-secondary-port
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure eaps failtime

```
configure eaps name failtime seconds milliseconds
```

Description

Configures the period after which the master node declares a failure if no hello PDUs are received.

Syntax Description

<i>name</i>	Specifies the name of an EAPS domain.
<i>seconds</i>	Specifies the number of seconds the master node waits before the failtimer expires. Default is 3 seconds, and the range is 0 to 300 seconds.
<i>milliseconds</i>	Specifies the number of milliseconds to wait before the failtimer expires. The range is 300 to 999 milliseconds.

Default

The 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 the failtimer expires. The failtime period (seconds plus milliseconds) must be set greater than the configured value for hellotime. The default value is three seconds.

Increasing the failtime value reduces the likelihood of false failure detections caused by network congestion.



Note

You configure the action taken when the failtimer expires by using the `configure eaps failtime expiry-action` command.

In ExtremeXOS 11.0, the failtimer range was 2 to 60 seconds.

Example

The following command configures the failtimer value for the EAPS domain eaps_1 to 15 seconds:

```
configure eapseaps_1failtime15 0
```

The following command configures the failtimer value for the EAPS domain eaps_2 to 300 milliseconds:

```
configure eapseaps_2failtime0 300
```

History

This command was first available in ExtremeXOS 11.0.

The range for the failtimer was changed to 2 to 300 seconds in ExtremeXOS 11.1. The default value for the failtimer remains unchanged.

The milliseconds parameter was added in ExtremeXOS 12.4.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure eaps fast-convergence

```
configure eaps fast-convergence[off | on]
```

Description

Enables EAPS to converge more quickly.

Syntax Description

off	Turns fast-convergence off. Default is off.
on	Turns fast-convergence on.

Default

Default is off.

Usage Guidelines

This command acts on the switch, not per domain.

In certain environments to keep packet loss to a minimum when the ring is broken, configure EAPS with fast-convergence turned on. If fast convergence is turned on, you can view the configuration with the [show eaps](#) command.



Note

If fast-convergence is turned on, the link filters on all EAPS ring ports are turned off. This can result problems if the port's hardware encountered a problem and started "flapping" between link-up/link-down states.

Example

The following command configures fast convergence for all of the EAPS domains on the switch:

```
configure eapsfast-convergence on
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure eaps hello-pdu-egress

```
configure eaps name hello-pdu-egress [primary-port | secondary-port]
```

Description

Configures the port through which a master node sends EAPS hello PDUs.

Syntax Description

<i>name</i>	Specifies the name of an EAPS domain.
-------------	---------------------------------------

Default

Default is the primary port.

Usage Guidelines

This command is provided for special network topologies that use spatial reuse and require that all EAPS hello PDUs travel in the same direction on the ring.



Note
We recommend the default (primary-port) configuration for this command.

Example

The following command configures the master switch to send EAPS hello packets from the secondary port:

```
configure eaps "domain12" hello-pdu-egress secondary-port
```

History

This command was first available in ExtremeXOS 12.4.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure eaps hellotime

```
configure eaps name hellotime seconds milliseconds
```

Description

Configures the period at which the master node sends EAPS hello PDUs to verify ring connectivity.

Syntax Description

<i>name</i>	Specifies the name of an EAPS domain.
<i>seconds</i>	Specifies the number of seconds to wait between transmission of hello PDUs on the control <u>VLAN</u> . The range is 0 to 15 seconds.
<i>milliseconds</i>	Specifies the number of milliseconds to wait between transmission of hello PDUs on the control VLAN. The range is 0 to 999 milliseconds.

Default

Default is 1 second.

Usage Guidelines

Use the `hellotime` keyword and its associated parameters to specify the amount of time the master node waits between transmissions of hello PDUs on the control VLAN. Increasing the `hellotime` value results in a reduced load on the processor and less traffic on the EAPS ring.



Note

The hello PDU timer value must be smaller than the fail timer value to prevent false failure detection. If you change the hello PDU timer, verify that the fail timer value remains larger.

This command applies only to the master node. If you configure the hello PDU timer for a transit node, the timer value is ignored. If you later reconfigure that transit node as the master node, the master node uses the configured hello PDU timer value.

In ExtremeXOS 11.0, the range is 1 to 15 seconds. If you are running ExtremeXOS 11.0 with the hello timer value greater than 15 seconds and you upgrade to ExtremeXOS 11.1 or later, you must modify the hello timer to be within the 1 to 15 seconds range.

Example

The following example configures the `hellotime` value for the EAPS domain `eaps_1` to 300 milliseconds:

```
configure eap seaps_1 hellotime 0 300
```

History

This command was first available in ExtremeXOS 11.0.

The range for the hello timer was changed to 1 to 15 seconds in ExtremeXOS 11.1. The default value for the hello timer remains unchanged.

Support for a specific number of milliseconds was added in ExtremeXOS 12.4.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure eaps mode

```
configure 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

<i>name</i>	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

One node (or switch) on the ring must be configured as the master node for the specified domain; all other nodes (or switches) on the ring are configured as transit nodes for the same domain.

If you configure a switch to be a transit node for an EAPS domain, the switch displays by default messages to:

- Remind you to configure a master node in the EAPS domain.
- Notify you that changing a master node to a transit node might cause a loop in the network. If you have not assigned a new master node before changing the current master node to a transit node, you might cause a loop in the network.

When prompted, do one of the following:

- Enter **y** to identify the switch as a transit node.
- Enter **n** or press **[Return]** to cancel this action.

If you have considerable knowledge and experience with EAPS, you might find the EAPS loop protection warning messages unnecessary. For more information, see the [configure eaps config-warnings off](#) command.

Example

The following example identifies this switch as the master node for the domain named eaps_1:

```
configure eaps eaps_1 mode master
```

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

```
configure eaps eaps_1 mode transit
```

The switch displays the following warning message and prompts you to confirm this action:

```
WARNING: Make sure this specific EAPS domain has a Master node in the ring. If you change
this node from EAPS master to EAPS transit, you could cause a loop in the network. Are you
sure you want to change mode to transit? (y/n)
```

Enter **y** to identify the switch as a transit node. Enter **n** to cancel this action.

History

This command was first available in ExtremeXOS 11.0.

The interactive messages were added in ExtremeXOS 11.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure eaps multicast add-ring-ports

```
configure eaps multicast add-ring-ports [on | off]
```

Description

Configures the switch to add previously blocked ring ports to existing multicast groups when an [EAPS](#) topology change occurs.

Syntax Description

on	Enables the multicast add-ring-ports feature.
off	Disables the multicast add-ring-ports feature.

Default

Off.

Usage Guidelines

When this feature is set to on and an EAPS topology change occurs, multicast traffic is fastpath forwarded using the switch hardware during the topology transition. The on setting improves multicast forwarding performance during the transition.



Note

EAPS multicast flooding must be enabled before this feature will operate. For information on enabling EAPS multicast flooding, see the [configure eaps multicast temporary-flooding](#) command description.

When this feature is set to off and an EAPS topology change occurs, multicast traffic is slowpath forwarded using the CPU during the topology transition. The off setting reduces multicast forwarding performance during the transition.

For other methods of supporting multicast traffic during an EAPS topology change, see the descriptions for the following commands:

- [configure eaps multicast send-igmp-query](#)
- [configure eaps multicast temporary-flooding](#)

Example

The following example enables the add-ring-ports feature:

```
configure eaps multicast add-ring-ports on
```

History

This command was first available in ExtremeXOS 12.1.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure eaps multicast send-igmp-query

```
configure eaps multicast send-igmp-query [on | off]
```

Description

Configures the switch to send *IGMP* query messages to all protected *VLANs* when an *EAPS* topology change occurs.

Syntax Description

on	Enables the multicast send-igmp-query feature.
off	Disables the multicast send-igmp-query feature.

Default

On.

Usage Guidelines

When this feature is set to on and an EAPS topology change occurs, the switch sends IGMP query messages to all protected VLANs. If the protected VLANs in the node detecting (and generating) the topology change do not have IP address, a query is generated with the source IP address set to the querier address in that VLAN.

In a EAPS ring with many protected VLANs, the many responses can impact switch performance. This is the default behavior and was the only method for supporting multicast traffic during EAPS topology changes prior to release 12.1.2.

When this feature is set to off and an EAPS topology change occurs, the switch does not automatically send IGMP queries to all protected VLANs during the topology transition. The off setting improves switch performance during the transition, but you should use one of the following commands to see that multicast traffic is supported during and after the topology change:

- `configure eaps multicast add-ring-ports`

- `configure eaps multicast temporary-flooding`

Example

The following command disables the send-igmp-query feature:

```
configure eaps multicast send-igmp-query off
```

History

This command was first available in ExtremeXOS 12.1.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure eaps multicast temporary-flooding duration

```
configure eaps multicast temporary-flooding duration seconds
```

Description

Configures the duration for which the switch temporarily enables multicast flooding when an *EAPS* topology change occurs.

Syntax Description

<i>seconds</i>	Specifies the period (in seconds) for which the switch enables multicast flooding.
----------------	--

Default

15 seconds.

Usage Guidelines

The flooding duration configuration applies only when the temporary-flooding feature is enabled with the following command:

```
configure eaps multicast temporary-flooding
```

Example

The following command configures the temporary-flooding feature duration for 30 seconds:

```
configure eaps multicast temporary-flooding duration 30
```

History

This command was first available in ExtremeXOS 12.1.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure eaps multicast temporary-flooding

```
configure eaps multicast temporary-flooding [on | off]
```

Description

Configures the switch to temporarily enable multicast flooding when an EAPS topology change occurs.

Syntax Description

on	Enables the multicast temporary-flooding feature.
off	Disables the multicast temporary-flooding feature.

Default

Off.

Usage Guidelines

When this feature is set to on and an EAPS topology change occurs, the switch temporarily enables multicast flooding to all protected VLANs for the duration specified by the following command:

```
configure eaps multicast temporary-flooding duration
```

If you change the configuration to off, topology changes that occur after this command do not result in temporary flooding. For example, if you change the configuration to off while flooding is in progress for a protected VLAN or set of protected VLANs (due to an EAPS topology change), the flooding continues for the configured duration period. New topology changes on the protected VLANs do not cause flooding.

When this feature is set to off and an EAPS topology change occurs, the switch does not enable flooding to all protected VLANs during the topology transition. The default switch response for multicast traffic during an EAPS topology change is that defined by the following command:

```
configure eaps multicast send-igmp-query
```

You can also use the following command to configure the switch response for multicast traffic during an EAPS topology change:

```
configure eaps multicast add-ring-ports
```


Example

The following command enables the temporary-flooding feature:

```
configure eaps multicast temporary-flooding on
```

History

This command was first available in ExtremeXOS 12.1.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure eaps name

```
configure eaps old_name name new_name
```

Description

Renames an existing *EAPS* domain.

Syntax Description

<i>old_name</i>	Specifies the current name of an EAPS domain.
<i>new_name</i>	Specifies a new name for the EAPS domain.

Default

N/A.

Usage Guidelines

If you use the same name across categories (for example, *STPD* and EAPS names), we recommend that you specify the identifying keyword as well as the actual name. If you do not use the keyword, the system might return an error message.

Example

The following command renames EAPS domain eaps-1 to eaps-5:

```
configure eaps eaps-1 name eaps-5
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure eaps port

```
configure eaps name [primary | secondary] port ports
```

Description

Configures a node port as the primary or secondary port for the specified *EAPS* domain.

Syntax Description

<i>name</i>	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.
<i>ports</i>	Specifies one port or slot and port.

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.

The primary and secondary ports have significance only on a master node. The health-check messages are sent out the primary port of the master node, and the master node blocks the protected *VLANs* on the secondary port.

The master node's secondary EAPS port cannot be configured on ports that are already configured as follows:

- Shared-port
- ISC port

There is no distinction between the primary and secondary ports on a transit node.

Beginning with ExtremeXOS 11.1, if you have a primary or secondary port that is a member of a load-shared group, you do not need to disable your EAPS domain and remove that ring port when modifying the load-shared group. For more information about configuring load sharing on your switch, see "Configuring Slots and Ports on a Switch" in the [Switch Engine 32.3 User Guide](#).

For complete information about software licensing, including how to obtain and upgrade your license and what licenses are appropriate for this feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

Messages Displayed when Adding EAPS Ring Ports to a VLAN

If you attempt to add EAPS ring ports to a VLAN that is not protected by EAPS, the switch prompts you by default to confirm this action. For example, if you use the `configure vlan vlan_name add ports port_list` command, and the ports that you are attempting to add to the VLAN are currently used by EAPS as either primary or secondary ring ports, the switch displays the following message:

```
Make sure <vlan_name> is protected by EAPS. Adding EAPS ring ports to a VLAN could cause a loop in the network. Do you really want to add these ports (y/n)
```

Enter y to add the ports to the VLAN. Enter n or press [Return] to cancel this action.

If you see this message, either configure the VLAN as an EAPS protected VLAN by using the `configure eaps add protected vlan` command or add ports that the EAPS domain does not use as primary or secondary ring ports.

If you have considerable knowledge and experience with EAPS, you might find the EAPS loop protection warning messages unnecessary. For more information, see the `configure eaps config-warnings off`.

Example

The following example adds port 1 to the EAPS domain eaps_1 as the primary port:

```
configure eapseaps_1primary port 1
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure eaps priority

```
configure eaps name priority {high | normal}
```

Description

Configures an *EAPS* domain priority.

Syntax Description

<i>name</i>	Specifies the name of an EAPS domain.
-------------	---------------------------------------

Default

Normal.

Usage Guidelines

Extreme Networks recommends that no more than 200 protected VLANs be configured as high priority domains. Priority protection works best when the majority of protected VLANs are configured for normal priority and a relatively small percentage of the protected VLANs are configured as high priority domains.

When EAPS domains on two separate physical rings share a common link (shared-port configuration) and have one or more protected VLANs in common, the domains must be configured with the same domain priority.

When EAPS domain priority is configured on separate physical rings that are connected to the same switch, the priorities on each ring are serviced independently. For example, if there is a break on both Ring A and Ring B, the high priority domains on each ring are serviced before the lower priority domains. However, the switch does not attempt to process the high priority domains on Ring B before servicing the normal priority domains on Ring A.

For a high priority domain to get priority over normal priority domains, all switches in the EAPS domain must support high priority domains. If high priority domains are configured on a switch that is in a ring with one or more switches that do not support high priority domains (software releases before ExtremeXOS Release 12.5), the high priority domain operates as a normal priority domain.

Example

The following command configures the eaps_1 domain as a high priority domain:

```
configure eapseaps_1 priority high
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure eaps shared-port common-path-timers

```
configure eaps shared-port port common-path-timers {[health-interval |  
timeout] seconds}
```

Description

Configures the common path health interval or timeout value.

Syntax Description

<i>port</i>	Specifies the port number of the common link port.
health-interval	Specifies the interval for health check messages on the common link.

timeout	Specifies the timeout value for the common link.
<i>seconds</i>	Specifies the amount of health interval, in seconds.

Default

N/A.

Usage Guidelines

This command allows you to configure the length of the common path health interval, in seconds, for a given port. The range is from 1 to 10 seconds.

Example

The following command configures a common-link health interval of 5 seconds on port 1:1.

```
configure eaps shared-port 1:1 common-path-timers health-interval 5
```

The following command configures a segment timeout of 10 seconds on port 1:1.

```
configure eaps shared-port 1:1 common-path-timers timeout 10
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms with the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and what licenses are appropriate for this feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure eaps shared-port link-id

```
configure eaps shared-port ports link-id id
```

Description

Configures the link ID of the shared port.

Syntax Description

<i>ports</i>	Specifies the port number of the common link port.
<i>id</i>	Specifies the link ID of the port. The link ID range is 1 to 65535.

Default

N/A.

Usage Guidelines

Each common link in the *EAPS* network must have a unique link ID. The controller and partner shared ports belonging to the same common link must have matching link IDs. No other instance in the network should have that link ID.

If you have multiple adjacent common links, we recommend that you configure the link IDs in ascending order of adjacency. For example, if you have an EAPS configuration with three adjacent common links, moving from left to right of the topology, configure the link IDs from the lowest to the highest value.

Example

The following command configures the EAPS shared port 1:1 to have a link ID of 1.

```
configure eaps shared-port 1:1 link-id 1
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms with the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and what licenses are appropriate for this feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure eaps shared-port mode

```
configure eaps shared-port ports mode controller | partner
```

Description

Configures the mode of the shared port.

Syntax Description

<i>ports</i>	Specifies the port number of the shared port.
<i>controller</i>	Specifies the controller mode. The controller is the end of the common link responsible for blocking ports when the common link fails thereby preventing the superloop.
<i>partner</i>	Specifies partner mode. The partner is responsible only for sending and receiving health-check messages.

Default

N/A.

Usage Guidelines

The shared port on one end of the common link must be configured to be the controller. This is the end responsible for blocking ports when the common link fails thereby preventing the superloop.

The shared port on the other end of the common link must be configured to be the partner. This end does not participate in any form of blocking. It is responsible only for sending and receiving health-check messages.

Example

The following command configures the shared port 1:1 to be the controller.

```
configure eaps shared-port 1:1 mode controller
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms with the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and what licenses are appropriate for this feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure eaps shared-port segment-timers expiry-action

```
configure eaps shared-port port segment-timers expiry-action [segment-down | send-alert]
```

Description

Configures the action taken when the segment timeout timer expires.

Syntax Description

<i>port</i>	Specifies the port number of the common link port.
segment-down	Marks the segment as DOWN if the segment timer expires. No link-status-query is sent to verify that links are down.
send-alert	If the segment timer expires, the switch keeps segments up, but sends a warning message to the log. The segment fail flag is set, an <i>SNMP</i> trap is sent, and a link-status-query is sent to verify if any links are down.

Default

Default is send-alert.

Usage Guidelines

By default, the action is to send an alert if the segment timeout timer expires. Instead of the segment going into a failed state and being marked as down, the segment remains in a segment up state with the failed flag set. The switch writes a critical error message to the syslog warning the user that there is a fault in the segment. An SNMP trap is also sent.



Note

Use caution when setting the segment-timeout expiry action to segment-down. Using this configuration, if the controller or partner node loses three consecutive hello PDUs, the failtimer expires—but there might not be a break in the segment. Opening a blocked port in this situation creates a loop.

The following describes some general recommendations for using this command:

- When you configure your Extreme Networks switches as the partner and controller, respectively, make sure that their segment timer configurations are identical.

For example, if you have a partner switch with the segment-timeout expiry action set to send-alert, make sure the controller switch has its segment-timeout expiry action set to send-alert.

However, if you have a partner switch with the segment-timeout expiry action set to send-alert, and the controller switch does not have a segment timer configuration, you must configure the partner switch's segment-timeout expiry action to segment-down.

- If you have a network containing non-Extreme Networks switches or non-*EAPS* devices, set the segment-timeout expiry action to segment-down.

The following events can cause a ring segment failure:

- There is a hardware failure.
- The controller or partner received a Link Down message from the partner or controller, respectively.
- The segment timer expires and the expiry action was set to segment-down. This means that either the controller or partner did not receive health check messages during the defined segment timeout period.

To view shared-port information, including shared-port segment status, use the following command:

```
show eaps shared-port {port}{detail}
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms with the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and what licenses are appropriate for this feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure eaps shared-port segment-timers health-interval

```
configure eaps shared-port port segment-timers health-interval seconds
```

Description

Configures the shared-port health interval timeout.

Syntax Description

<i>port</i>	Specifies the port number of the common link port.
<i>seconds</i>	Specifies the amount of health interval, in seconds.

Default

N/A.

Usage Guidelines

This command allows you to configure the length of the shared-port health interval timeout, in seconds, for a given port.

Example

The following command configures a shared-port health interval timeout of 10 seconds on port 1:1.

```
configure eaps shared-port 1:1 segment-timers health-interval 10
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms with the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and what licenses are appropriate for this feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure eaps shared-port segment-timers timeout

```
configure eaps shared-port port segment-timers timeout seconds
```

Description

Configures the shared-port timeout.

Syntax Description

<i>port</i>	Specifies the port number of the common link port.
<i>seconds</i>	Specifies the amount of health interval, in seconds.

Default

N/A.

Usage Guidelines

This command allows you to configure the length of the shared-port timeout, in seconds, for a given port.

Example

The following command configures a shared-port timeout of 10 seconds on port 1:1.

```
configure eaps shared-port 1:1 segment-timers timeout 10
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms with the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and what licenses are appropriate for this feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure edp advertisement-interval

```
configure edp advertisement-interval timer holddown-interval timeout
```

Description

Sets the advertisement interval and hold down interval for *EDP*.

Syntax Description

<i>timer</i>	Specifies the advertisement interval in seconds.
<i>timeout</i>	Specifies the hold down interval in seconds.

Default

The default setting for timer is 60 seconds, and for timeout is 180 seconds.

Usage Guidelines

Extreme Discover Protocol (EDP) is used to gather information about neighbor Extreme Networks switches. EDP-enabled ports advertise information about the Extreme switch to other switches on the interface and receive advertisements from other Extreme switches. Information about other Extreme switches is discarded after the hold down interval timeout value is reached without receiving another advertisement.

Example

The following command configures the EDP advertisement-interval to 2 minutes and the hold down interval to 6 minutes:

```
configure edp advertisement-interval 120 holddown-interval 360
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure elrp-client dynamic-vlans

```
configure elrp-client dynamic-vlans {mvrp | netlogin | vm-tracking |  
  policy | fabric-attach | all} [on | off]
```

Description

This command enables/disables Extreme Loop Recognition Protocol (ELRP) over various types of dynamic VLANs.

Syntax Description

dynamic-vlans	ELRP configuration options for dynamically created VLANs.
mvrp	Specifies that the command applies to dynamic VLANs created by Multiple VLAN Registration Protocol (MVRP) only.
netlogin	Specifies that the command applies to dynamic VLANs created by Network Login only.
vm-tracking	Specifies that the command applies to dynamic VLANs created by virtual machine MAC tracking only.

policy	Specifies that the command applies to dynamic VLANs created by One Policy.
fabric-attach	Specifies that the command applies to dynamic VLANs created by Fabric Attach.
all	(Default) Specifies that the command applies to all types of dynamic VLANs.
on	Enable ELRP for dynamic VLANs.
off	Disables ELRP for dynamic VLANs.

Default

ELRP for dynamic VLANs is "off" by default. If the type of dynamic VLAN is not specified, the command applies to all types of dynamic VLANs.

Example

The following example enables ELRP for all types of dynamic VLANs:

```
configure elrp-client dynamic-vlans on
```

The following example disables ELRP for VM tracking dynamic VLANs:

```
configure elrp-client dynamic-vlans vm-tracking off
```

The following example enables ELRP for Netlogin dynamic VLANs:

```
configure elrp-client dynamic-vlans netlogin on
```

The following example enables ELRP for One Policy dynamic VLANs:

```
configure elrp-client dynamic-vlans policy on
```

The following example disables ELRP for One Policy dynamic VLANs:

```
configure elrp-client dynamic-vlans policy off
```

History

This command was first available in ExtremeXOS 22.2.

The **policy** and **fabric-attach keywords** were added in ExtremeXOS 32.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure elrp-client dynamic-vlans action

```
configure elrp-client dynamic-vlans {mvrp | netlogin | vm-tracking |  
    policy | fabric-attach | all} [{interval sec} {action [{log | trap |
```

```
log-and-trap { disable-port { [{ egress | ingress } { duration seconds | permanent } ] | none } ] }
```

Description

This command sets actions to be taken after Extreme Loop Recognition Protocol (ELRP) on dynamic VLANs detects a loop.

Syntax Description

dynamic-vlans	ELRP configuration options for dynamically created VLANs.
mvrp	Specifies that the command applies to dynamic VLANs created by Multiple VLAN Registration Protocol (MVRP) only.
netlogin	Specifies that the command applies to dynamic VLANs created by Network Login only.
vm-tracking	Specifies that the command applies to dynamic VLANs created by virtual machine MAC tracking only.
policy	Specifies that the command applies to dynamic VLANs created by One Policy.
fabric-attach	Specifies that the command applies to dynamic VLANs created by Fabric Attach.
all	(Default) Specifies that the command applies to all types of dynamic VLANs.
interval	Specifies setting the time interval between successive ELRP polls.
<i>sec</i>	Sets the time interval in seconds between successive ELRP polls. Range is 1-600. Default = 1.
action	Action to be taken after ELRP poll result.
log	Print ELRP poll result to system log.
trap	Send SNMP trap.
log-and-trap	Print ELRP poll result to system log and send SNMP trap.
disable-port	Disable port where looped PDU was transmitted or received.
egress	Disable port where looped PDU was transmitted.
ingress	Disable port where looped PDU was received (default).
duration	Specifies setting the time period that the port is kept disabled before re-enabling.
<i>seconds</i>	Sets the time in seconds that the port is kept disabled before re-enabling. Range is 15-600. Default = 30.
permanent	Keep port disabled permanently. You must intervene to re-enable.
none	Removes any previously set actions.

Default

If the type of dynamic VLAN is not specified, the command applies to all types of dynamic VLANs.

If the time duration is not set for the period between ELRP polls, the default is one second.

If not specified, the port that the looped PDU was received on is disabled.

If not specified, the disabled port is kept disabled for 30 seconds before it is re-enabled.

Example

The following example enables ELRP for all types of dynamic VLANs with a time interval between ELRP polls of 2 seconds:

```
configure elrp-client dynamic-vlans interval 2
```

The following example enables ELRP for MVRP VLANs with SNMP trap set when a loop is detected:

```
configure elrp-client dynamic-vlans mvrp action trap
```

The following example enables ELRP for all types of dynamic VLANs and disables the egress port where the loop is detected permanently:

```
configure elrp-client dynamic-vlans action disable-port egress permanent
```

The following example enables ELRP for VM-tracking VLANs and disables the ingress port where the loop is detected for 100 seconds:

```
configure elrp-client dynamic-vlans vm-tracking action disable-port ingress duration 100
```

The following example enables ELRP for Fabric Attach VLANs and disables the ingress port where the loop is detected for 100 seconds:

```
configure elrp-client dynamic-vlans fabric-attach action disable-port ingress duration 100
```

History

This command was first available in ExtremeXOS 22.3.

The **policy** and **fabric-attach keywords** were added in ExtremeXOS 31.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure elrp-client dynamic-vlans client/uplink ports/remote-endpoints vxlan

```
configure elrp-client dynamic-vlans [netlogin | vm-tracking] [client-  
ports | uplink-ports | remote-endpoints vxlan] [on | off]
```

Description

This command turns Extreme Loop Recognition Protocol (ELRP) on/off for client ports or uplink ports for dynamic VLANs.

Syntax Description

netlogin	Specifies that the command applies to dynamic VLANs created by Network Login only.
vm-tracking	Specifies that the command applies to dynamic VLANs created by virtual machine MAC tracking only.
client-ports	Specifies client ports only.
uplink-ports	Specifies uplink ports only.
remote-endpoints	Specifies remote endpoints that are part of this VLAN.
vxlan	Specifies VXLAN remote endpoints that are part of this VLAN.
on	Enable ELRP for dynamic VLANs.
off	Disables ELRP for dynamic VLANs.

Default

ELRP for dynamic VLANs is "off" by default.

Example

The following example enables ELRP for Netlogin dynamic VLANs on uplink ports only:

```
configure elrp-client dynamic-vlans netlogin uplink-ports on
```

The following example enables ELRP for XNV dynamic VLANs on VXLAN remote endpoints that belong to the VLAN:

```
configure elrp-client dynamic-vlans vm-tracking remote-endpoints vxlan on
```

History

This command was first available in ExtremeXOS 22.2.

Remote endpoint capability was added in ExtremeXOS 22.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure elrp-client disable ports

```
configure elrp-client disable-ports [exclude | include] [ ports | eaps-  
ring-ports | remote-endpoints vxlan | inter-vlan-loop]
```

Description

Creates an ELRP exclude port list.

Syntax Description

exclude	Specifies that selected ports are to be excluded from ELRP disabling.
include	Specifies that selected ports are to be included in ELRP disabling.
<i>ports</i>	Specifies one or more ports to be excluded or included.
eaps-ring-ports	Specifies whether <i>EAPS</i> ring ports are to be excluded or included.
remote-endpoints	Specifies remote endpoints, if any, that are part of this VLAN.
vxlan	Specifies VXLAN remote endpoints that are part of this VLAN.
inter-vlan-loop	Excludes inter-VLAN loop detected ports.

Default

All ports, together with EAPS ring ports and VXLAN remote endpoints, are included by default; that is, they are disabled if a loop is detected on that port.

Usage Guidelines

Use this command to specify ports, EAPS ring ports, or VXLAN remote endpoints that are to be part of an ELRP exclude port list. Use the `exclude` option to add ports to the exclude port list. Use the `include` option to remove them from the list.

When ELRP detects a loop and has been configured to automatically disable the port where a looped ELRP PDU is received and an exclude port list has been configured, it will check to determine if that port is on the exclude port list. If that port is on the list, ELRP will not disable it; if it is not on the list, it will be disabled.

To display the ports that are include in the exclude port list, use the `show elrp disabled-ports` command.

To remove the exclude port list, use the `unconfigure elrp-client disable ports` command.

Example

The following example adds port 2:1 to an ELRP exclude port list:

```
configure elrp-client disable-ports exclude 2:1,2:3
```

History

This command was first available in ExtremeXOS 12.5.3.

VXLAN remote endpoint option added in ExtremeXOS 22.4.

The **inter-vlan-loop** option for excluding inter-VLAN loop detected ports was added in ExtremeXOS 30.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure elrp-client hardware-assist loopback-port

```
configure elrp-client hardware-assist loopback-port [port | none]
```

Description

Configures or unconfigures a front panel port as the designated loopback port for hardware-assisted *ELRP (Extreme Loop Recovery Protocol)*.

Syntax Description

elrp-client	Configures ELRP client.
hardware-assist	Selects configuring hardware-assisted ELRP.
loopback-port	Designates selecting a loopback port for hardware-assisted ELRP.
<i>port</i>	Selects the loopback port. The port must be an unused front panel port.
none	Unconfigures a loopback port.

Default

N/A.

Usage Guidelines

The loopback port must be an unused front panel port. The selected loopback port cannot be part of a VLAN. The loopback port cannot be changed or unconfigured if hardware-assisted ELRP mode is enabled. To disable hardware-assisted ELRP, use the command `disable elrp-client`.

Example

The following example configures port 7 as the loopback port for hardware-assisted ELRP:

```
# configure elrp-client hardware-assist loopback-port 7
```

The following example unconfigures the loopback port for hardware-assisted ELRP:

```
# configure elrp-client hardware-assist loopback-port none
```

History

This command was first available in ExtremeXOS 30.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure elrp-client inter-vlan-loop-detection

```
configure elrp-client inter-vlan-loop-detection [on | off]
```

Description

Turns on/off Extreme Loop Recovery Protocol (ELRP) inter-VLAN loop detection.

Syntax Description

inter-vlan-loop-detection	ELRP detects loops between untagged ports on different VLANs on the same switch.
on	Turns on Inter-VLAN loop detection. (Default)
off	Turns off Inter-VLAN loop detection.

Default

Inter-VLAN loop detection is on by default.

Usage Guidelines

It is common in networks for you to accidentally inter-connect two different VLANs by looping together two untagged ports (one in each respective VLAN). This type of configuration results in an outage, and it is difficult for the average user to detect.

If desired, you can then include or exclude the inter-VLAN loops to be disabled using the `configure elrp-client disable-ports [exclude | include] [ports | eaps-ring-ports | remote-endpoints vxlan | inter-vlan-loop]` command.

Example

The following example turns on Inter-VLAN loop detection:

```
# configure elrp-client inter-vlan-loop-detection on
```

History

This command was first available in ExtremeXOS 30.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure elrp-client one-shot

```
configure elrp-client one-shot {vlan [vlan_name | all] ports [ports |
    all | none] {remote-endpoints vxlan all} {interval interval {seconds |
    milliseconds}} {retry count} {log | print | print-and-log}}
```

Description

Starts one-time, non-periodic ELRP packet transmission on the specified ports of the VLAN using the specified count and interval.

Syntax Description

<i>vlan_name</i>	Specifies a VLAN name.
all	Specifies all VLANs.
<i>ports</i>	Specifies the set of VLAN ports for packet transmission.
remote-endpoints	Specifies remote endpoints that are part of this VLAN.
vxlan	Specifies VXLAN remote endpoints that are part of this VLAN.
interval	Time interval between two successive ELRP PDUs.
<i>interval</i>	Interval value between 1–64 seconds or 100–64,000 milliseconds. Default is 1 second.
seconds	Specifies that time interval is in the unit of seconds.
milliseconds	Specifies that time interval is in the unit of milliseconds.
all	Specifies all ports of this VLAN for packet transmission.
<i>count</i>	Specifies the number of times ELRP packets must be transmitted. The range is 1 to 255 times. The default is 3 times.
log	Specifies that a message should be logged in the system log file when ELRP packets are received back indicating detection of network loop, or no packets are received within the specified duration.
print	Specifies that a message should be printed to the console when ELRP packets are received back indicating detection of network loop, or no packets are received within the specified duration.
print-and-log	Specifies that a message should be logged in the system log file and printed to the console when ELRP packets are received back indicating detection of network loop, or no packets are received within the specified duration.

Default

Second—The interval between consecutive packet transmissions is 1 second.

Count—The number of time ELRP packets must be transmitted is 10.

Usage Guidelines

This command starts one-time, non-periodic ELRP packet transmission on the specified ports of the VLAN using the specified count and interval. If any of these transmitted packets is returned, indicating loopback detection, the ELRP client can perform a configured action such as logging a message in the system log file or printing a log message to the console. There is no need to send a trap to the *SNMP* manager for non-periodic requests.



Note

You can also use the command [run elrp](#) on page 2425 to perform one-time ELRP packet transmission.

Use the [configure elrp-client periodic](#) command to configure periodic transmission of ELRP packets.

The ELRP client must be enabled globally in order for it to work on any VLANs. Use the [enable elrp-client](#) command to globally enable the ELRP client.

The ELRP client can be disabled globally so that none of the ELRP VLAN configurations take effect. Use the [disable elrp-client](#) command to globally disable the ELRP client.

Example

The following example starts one-time, non-periodic ELRP packet transmission on all ports of the VLAN sales, uses the default interval and transmission times, and sends messages to the console:

```
configure elrp-client one-shot sales ports all interval 1 seconds retry 3 print
```

History

This command was first available in ExtremeXOS 11.1.

The ability to specify the time interval in milliseconds was introduced in ExtremeXOS 22.4.

VXLAN remote endpoint option added in ExtremeXOS 22.4.

The **all** option for VLANs was added in ExtremeXOS 30.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure elrp-client periodic

```
configure elrp-client periodic {vlan} vlan_name ports [ports | all |
none] {remote-endpoints vxlan all} {interval interval {seconds |
milliseconds}} {log | log-and-trap | trap} {disable-port {egress |
ingress} {duration {seconds} | permanent}}
```

Description

Starts periodic ELRP packet transmission on the specified ports of the VLAN or VXLAN remote tunnel endpoints (RTEPs) using the specified interval.

Syntax Description

vlan	Specifies a VLAN name.
<i>vlan_name</i>	Specifies a VLAN name.
<i>ports</i>	Specifies the set of VLAN ports for packet transmission.
all	Specifies all ports for packet transmission.
none	Specifies no ports for packet transmission. This option allows you to configure (unambiguously) ELRP on only VXLAN RTEPs.
remote-endpoints	Specifies to include the remote endpoints, if any, in this VLAN. Only supported with software ELRP.
vxlan	Specifies VXLAN remote endpoints.
interval	Time interval between two successive ELRP PDUs.
<i>interval</i>	Software ELRP interval range between 1–600 seconds or 1,000–600,000 ms. Hardware-assisted ELRP interval range is between 3–600,000 ms. Default is 1 second.
seconds	Specifies that time interval is in the unit of seconds.
milliseconds	Specifies that time interval is in the unit of milliseconds.
log	Specifies that a message should be logged in the system log file when ELRP packets are received back indicating detection of network loop, or no packets are received within the specified duration.
log-and-trap	(Default) Specifies that a message should be logged in the system log file and trap message should be sent to the <u>SNMP</u> manager when ELRP packets are received back indicating detection of network loop, or no packets are received within the specified duration.
trap	Specifies that a trap message should be sent to the SNMP manager when ELRP packets are received back indicating detection of network loop, or no packets are received within the specified duration.
disable-port	Specifies that the port should be disabled where the looped PDU is received.
egress	Disable port where looped PDU was transmitted. Only supported with software ELRP.
duration	Specifies a hold time that the port is kept disabled before re-enabling.
<i>seconds</i>	The number of seconds the port is kept disabled.
permanent	Specifies that the port is disabled permanently. User intervention is required to enable.

Default

The default interval between consecutive packet transmissions is 1 second.

If a duration in seconds is not specified, the default is permanent.

If not specified, log-and-trap action is the default.

Usage Guidelines

This command starts periodic ELRP packet transmission on the specified ports of the VLAN using the specified interval. If any of these transmitted packets is returned, indicating loopback detection, the ELRP client performs a configured action of logging a message in the system log file and/or sending a trap to the SNMP manager.

Beginning with ExtremeXOS 12.4, you have the option to automatically disable the port where the looped packet arrives and to specify the time interval for which the port remains disabled. When that specified time expires, the port is automatically enabled.

Should a loop occur on multiple ports, only the first port in the VLAN on which the PDU is received is disabled. The second port is ignored for 1 or 2 seconds and then if another PDU is received, that port is disabled until the loop is gone. This prevents shutting down all ports in the VLAN.

Use either the `configure elrp-client one-shot` or the `run elrp` command to configure non-periodic, one-time transmission of ELRP packets.

The ELRP client must be enabled globally in order for it to work on any VLANs. Use the `enable elrp-client` command to globally enable the ELRP client.

The ELRP client can be disabled globally so that none of the ELRP VLAN configurations take effect. Use the `disable elrp-client` command to globally disable the ELRP client.

Use the `show elrp` command to check the ELRP status and the `show elrp disabled-ports` command to view details of ELRP disabled ports.

For the **interval** option with hardware-assisted ELRP, hardware-assisted ELRP uses ACL meter to rate limit the PDU TX rate, which has a granularity of 8 Kbps, so for any interval configured for longer than 70ms in hardware-assisted ELRP mode, the actual interval is around 70ms. This is determined by hardware capabilities of the switch.

Example

The following example starts periodic ELRP packet transmission on slot 3, port 2 of VLAN marketing, sends packet transmissions every 2 seconds, sends messages to the log, and should a loop be detected, disables the port for 5 seconds:

```
configure elrp-client periodic marketing ports 3:2 interval 2 seconds log disable-port
duration 5
```

History

This command was first available in ExtremeXOS 11.1.

The disable port feature was added in ExtremeXOS 12.4.

The ability to specify the time interval in milliseconds was introduced in ExtremeXOS 22.4.

The ability to specify VXLAN RTEPs was introduced in ExtremeXOS 22.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure elsm ports hellotime

```
configure elsm ports port_list hellotime hello_time
```

Description

Configures the ELSM hello timer by specifying the time between consecutive hello messages for the specified ports.

Syntax Description

<i>port_list</i>	Specifies the port or ports for which the ELSM hello timer should be configured.
<i>hello_time</i>	Specifies the time in seconds between consecutive hello messages. Use the same value for the hello interval on peer ports. The default value is 1 second, and the range is 1 to 128 seconds.

Default

The default is 1 second.

Usage Guidelines

ELSM works between two connected ports, and each ELSM instance is based on a single port.

When you enable ELSM on the specified ports, the ports participate in ELSM with their peers and begin exchanging ELSM hello messages.

ELSM uses two types of hello messages to communicate the health of the network to other ELSM ports:

- Hello+ — The ELSM-enabled port receives a hello message from its peer and no problem is detected.
- Hello- — The ELSM-enabled port does not receive a hello message from its peer.

ELSM also has hello transmit states. The hello transmit states display the current state of transmitted ELSM hello messages. For more information about the hello transmit states, see the [show elsm ports](#) command.

A high hello timer value can increase the time it takes for the ELSM-enabled port to enter the Up state. The down timer is (2 + hold threshold) * hello timer. Assuming the default value of 2 for the hold threshold, configuring a hello timer of 128 seconds creates a down timer of (2 + 2) 128, or 512 seconds. In this scenario it would take 512 seconds for the port to transition from the Down to the Up state.

If you modify the hello timer on one port, we recommend that you use the same hello timer value on its peer port.

Example

The following command specifies 5 seconds between consecutive ELSM hello messages for slot 2, ports 1-2 on the switch:

```
configure elsm ports 2:1-2:2 hellotime 5
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure elsm ports hold-threshold

```
configure elsm ports port_list hold-threshold hold_threshold
```

Description

Configures the number of Hello+ messages required by the specified ELSM-enabled ports to transition from the Down-Wait state to the Up state.

Syntax Description

<i>port_list</i>	Specifies the port or ports for which the ELSM hold threshold should be configured.
<i>hold_threshold</i>	Specifies the number of Hello+ messages required to transition from the Down-Wait state to the Up state. The default is 2 messages, and the range is 1 to 40 messages.

Default

The default is 2 Hello+ messages.

Usage Guidelines

The port begins in the Down state, so the first received Hello+ message transitions the ELSM-enabled port from the Down state to the Down-Wait state. After that transition, the configured hold-threshold value determines the number of Hello+ messages required to transition from Down-Wait state to the Up state.

The ELSM hold threshold determines the number of Hello+ messages the ELSM peer port must receive to transition from the Down-Wait state to the Up state. For example, a threshold of 1 means the ELSM port must receive at least one Hello+ message to transition from the Down-Wait state to the Up state.

After the down timer expires, the port checks the number of Hello+ messages against the hold threshold. If the number of Hello+ messages received is greater than or equal to the configured hold threshold, the ELSM receive port moves from the Down-Wait state to the Up state.

If the number of Hello+ messages received is less than the configured hold threshold, the ELSM receive port moves from the Down-Wait state back to the Down state and begins the process again.

If you modify the hold threshold on one port, we recommend that you use the same hold threshold value on its peer port.

You configure the hold threshold on a per-port basis, not on a per-switch basis.

Example

The following command specifies that two Hello+ messages are required for the ELSM receive ports configured on slot 2, ports 1-2, to transition from the Down-Wait state to the Up state:

```
configure elsm hold-threshold 2 ports 2:1-2:2
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure elsm ports *port_list* **uptimer-threshold** *uptimer_threshold*

```
configure elsm ports port_list uptimer-threshold uptimer_threshold
```

Description

Configures the number of Hello+ messages required by the specified ELSM-enabled ports to transition from the Up state to the Down state.

Syntax Description

<i>port_list</i>	Specifies the port or ports for which the ELSM hold threshold should be configured.
<i>uptimer_threshold</i>	Specifies the number of Hello+ messages required to transition from the Up- state to the Down state. The default is 6messages, and the range is 3 to 60 messages.

Default

The default is 6 Hello+ messages.

Usage Guidelines

The ELSM up timer begins when the ELSM-enabled port enters the UP state. Each time the port receives a Hello+ message, the timer restarts. Up timer is `Uptimer_threshold * hello timer`. When the Up timer expires, it transits from UP state to DOWN state.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure erps add control vlan

```
configure erps ring-name add control {vlan} vlan_name
```

Description

Add a control VLAN on the ERPS ring.

Syntax Description

<i>ring-name</i>	Alphanumeric string that identifies the ERPS ring.
control	VLAN that carries ERPS control traffic.
<i>vlan_name</i>	Alphanumeric string identifying the VLAN to be used for control traffic.

Default

N/A.

Usage Guidelines

Use this command to add a control VLAN on the ERPS ring. This is the VLAN that carries ERPS control traffic.



Note

Other VLAN types such as VMAN, SVLAN, CVLAN and BVLAN will not be used for control traffic. A control VLAN cannot be deleted from a ring that has CFM configured.

Example

The following command adds a control VLAN named “vlan10” to an ERPS ring named “ring1”:

```
configure erps ring1 add control vlan vlan10
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

configure erps add protected vlan

```
configure erps ring-name add protected {vlan} vlan_name
```

Description

Add a protected VLAN on the ERPS ring. This is a data VLAN that ERPS will protect.

Syntax Description

<i>ring-name</i>	Alphanumeric string that identifies the ERPS ring.
<i>vlan_name</i>	Alphanumeric string identifying the data VLAN to be added that ERPS will protect. This can be a VLAN, SVLAN, BVLAN or VMAN.

Default

N/A.

Usage Guidelines

Use this command to add a protected data VLAN on the ERPS ring. This VLAN will be protected by ERPS, and it can be a VLAN, SVLAN, BVLAN or VMAN.



Note

The SVLAN-BVLAN combination cannot both be added to the same ring or sub-ring.

Example

The following command adds a protected VLAN named “vlan10” to an ERPS ring named “ring1”:

```
configure erps ring1 add protected vlan vlan10
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

configure erps control-mac

```
configure erps ring-name control-mac [auto | default]
```

Description

Configures *ERPS* control MAC (either default or auto) on a particular ERPS ring instance.

Syntax Description

erps	Specifies ERPS (ITU-T G.8032).
<i>ring-name</i>	Specifies the alphanumeric string that identifies the ERPS ring/sub-ring.
control-mac	Destination MAC used in R-APS PDUs .
auto	Use ring ID-based MAC address (01:19:A7:00:00:ring-id).
default	Use default MAC address (01:19:A7:00:00:01).

Default

By default, if an ERPS ring instance is created with a user-defined ring ID, the control MAC used by ring instance is auto (01:19:A7:00:00:ring-id).

By default, if an ERPS ring instance is created without a user-defined ring ID, the control MAC used by the ring instance is default (01:19:A7:00:00:01).

Usage Guidelines

As per the ITU G.8032 standard, destination MAC used in R-APS PDUs are of 2 types:

- 01:19:A7:00:00:01 (default)
- 01:19:A7:00:00:ringId (auto)



Note

This command is only applicable on ERPS ring instances created with user-defined ring ID.

Example

The following example configures the control MAC of an ERPS ring instance created with a user-defined ring ID:

```
# configure erps Ring2 control-mac auto
```

The following example configures the control MAC of an ERPS ring instance created without a user-defined ring ID:

```
# configure erps Ring1 control-mac auto
    Error: This cli is applicable only when the erps ring is created with a user
    defined ringId. The default
    control-mac is used here.
```

History

This command was first available in ExtremeXOS 32.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure erps cfm port group

```
configure erps ring_name cfm port [east | west] [add | delete] group
group_name
```

Description

Associates or disassociates fault monitoring entities on the ERPS ring ports.

Syntax Description

<i>ring_name</i>	Alphanumeric string that identifies the ERPS ring.
east	East port.
west	West port.
add	Associates a CFM Down-MEP entity.
delete	Disassociates a CFM Down-MEP entity.
group	Specifies a CFM Down-MEP group.
<i>group_name</i>	Specifies the name of the Down MEP group.

Default

N/A.

Usage Guidelines

Use this command to associate or disassociate fault monitoring entities on the ERPS ring ports.

Example

The following command associates fault monitoring on the group "group1":

```
configure erps ring1 cfm port east add group1
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on all platforms running ExtremeXOS.

configure erps cfm protection group

```
configure erps ring_name cfm protection [add delete] group cfm_group
```

Description

Associates or disassociates a CFM UP MEP group for subring protection across the main ring.

Syntax Description

<i>ring-name</i>	Alphanumeric string that identifies the <i>ERPS</i> ring.
add	Associates a CFM Up-MEP entity.
delete	Disassociates a CFM Up-MEP entity.
group	Specifies a CFM Up-MEP group.

Default

N/A.

Usage Guidelines

Use this command to associate or disassociate a CFM UP MEP group for subring protection across the main ring.

When an UP MEP is configured for protection of a subring, the Manual Switch event will be enforced on the subring port on the interconnected nodes. As per Appendix X of the standard, the MS is issued when the node type and the multiple failure type are the same. ExtremeXOS implementation currently configures the node type to be the same as the fault type. So the user will notice both the

subring ports of the two interconnected nodes to be held in MS when multiple failures on the main ring occur. When the multiple failure clears this MS is also cleared.

Example

The following command associates a CFM UP MEP group for subring protection on the group "group1":

```
configure erps ring1 cfm protection add group1
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on all platforms running ExtremeXOS.

configure erps delete control vlan

```
configure erps ring-name delete control {vlan} vlan_name
```

Description

Delete a control VLAN on the ERPS ring.

Syntax Description

<i>ring-name</i>	Alphanumeric string that identifies the ERPS ring.
<i>vlan_name</i>	Alphanumeric string identifying the VLAN used for control traffic.

Default

N/A.

Usage Guidelines

Use this command to delete a control VLAN from the ERPS ring. This is the VLAN that carries ERPS control traffic.



Note

Other VLAN types such as VMAN, SVLAN, CVLAN and BVLAN will not be used for control traffic.

A control VLAN cannot be deleted from a ring that has CFM configured.

Example

The following command deletes a control VLAN named “vlan10” from an ERPS ring named “ring1”:

```
configure erps ring1 delete control vlan vlan10
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

configure erps delete protected vlan

```
configure erps ring-name delete protected {vlan} vlan_name
```

Description

Delete a protected data *VLAN* from the *ERPS* ring.

Syntax Description

<i>ring-name</i>	Alphanumeric string that identifies the ERPS ring.
<i>vlan_name</i>	Alphanumeric string identifying the data VLAN to be deleted from the ERPS ring.

Default

N/A.

Usage Guidelines

Use this command to delete a protected VLAN from the ERPS ring.

Example

The following command deletes a protected VLAN named “vlan10” from an ERPS ring named “ring1”:

```
configure erps ring1 delete protected vlan vlan10
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

configure erps dynamic-state

```
configure erps ring-name dynamic-state [force-switch | manual-switch |  
clear] port slot:port
```

Description

Configure or clear force-switch or manual-switch for the *ERPS* ring/sub-ring.

Syntax Description

dynamic-state	Configure force/manual/clear switch on the active ERPS ring.
force-switch	Force-switch operation.
manual-switch	Manual-switch operation.
clear	Clears force-switch/manual-switch.

Default

N/A.

Usage Guidelines

Use this command to configure or clear force-switch or manual-switch for the ERPS ring/sub-ring.



Note

In non-revertive mode, in the "Pending" state, you can use the **clear** option of this command to return to the "Idle" state where the blocked link is manually reverted to the Ring Protection Link (RPL).

Example

The following command clears force-switch and manual-switch on an ERPS ring named "ring1":

```
configure erps ring1 dynamic-state clear
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure erps name

```
configure erps old-ring-name name new-ring-name
```

Description

Rename the *ERPS* ring/sub-ring.

Syntax Description

<i>old-ring-name</i>	Alphanumeric string that identifies the ERPS ring.
<i>new-ring-name</i>	New alphanumeric string identifying the ERPS ring.

Default

N/A.

Usage Guidelines

Use this command to rename the ERPS ring or sub-ring.

Example

The following command an ERPS ring from “ring1” to “ring2”:

```
configure erps ring1 name ring2
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

configure erps neighbor port

```
configure erps ring-name neighbor-port port
```

Description

Add RPL (ring protection link) neighbor configuration for the *ERPS* ring.

Syntax Description

<i>ring-name</i>	Alphanumeric string that identifies the ERPS ring.
<i>port</i>	The slot:port number for RPL neighbor.

Default

N/A.

Usage Guidelines

Use this command to add RPL neighbor configuration for the ERPS ring.



Note

This command implicitly makes the node on which it is configured the RPL neighbor.

Example

The following command adds RPL neighbor on port 5 to an ERPS ring named “ring1”:

```
configure erps ring1 neighbor-port 5
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

configure erps notify-topology-change

```
configure {erps} ring-name notify-topology-change {eaps} domain_name
```

Description

Add an ERPS sub-ring to the EAPS domain.

Syntax Description

<i>ring-name</i>	Alphanumeric string identifying the ERPS sub-ring.
<i>domain_name</i>	Alphanumeric string identifying the EAPS domain.

Default

N/A.

Usage Guidelines

Use this command to add an ERPS sub-ring to the EAPS domain.

Example

Example output not yet available and will be provided in a future release.

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

configure erps protection-port

```
configure erps ring-name protection-port port
```

Description

Add ring protection link (RPL) owner configuration for the *ERPS* ring.

Syntax Description

<i>ring-name</i>	Alphanumeric string that identifies the ERPS ring.
<i>port</i>	The slot:port number for the ring protection link (RPL) owner.

Default

N/A.

Usage Guidelines

Use this command to add ring protection link (RPL) owner configuration for the ERPS ring.



Note

This command implicitly makes the node on which it is configured the RPL owner.

Example

The following command adds RPL owner configuration on port 5 to an ERPS ring named “ring1”:

```
configure erps ring1 protection-port 5
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

configure erps revert

```
configure {erps} ring-name revert [ enable | disable ]
```

Description

Add or delete *ERPS* revert operation along with the “wait-to-restore” time interval.

Syntax Description

<i>ring-name</i>	Alphanumeric string that identifies the ERPS ring.
enable	Enable revert mode to ERPS ring.
disable	Disable revert mode from ERPS ring.

Default

The default is the revertive mode (enable).

Usage Guidelines

Use this command to enable/disable a G.8032 ring to revert to the original ring protection link (RPL) block state.

Example

The following command disables revert mode from an ERPS ring named “ring1”:

```
configure erps ring1 revert disable
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

configure erps ring-ports east | west

```
configure erps ring-name ring-ports [east | west] port
```

Description

Add ring ports on the *ERPS* ring. The ring ports connect the switch to the ERPS ring.

Syntax Description

<i>ring-name</i>	Alphanumeric string that identifies the ERPS ring.
east	Add the ring port to the east port of the switch.
west	Add the ring port to the west port of the switch.
<i>port</i>	The slot:port number for the ring port.

Default

N/A.

Usage Guidelines

Use this command to add ring ports on the ERPS ring. The ring ports can be added to the east or west port of the switch. The ring ports connect the switch to the ERPS ring.

Example

The following command adds port 5 as a ring port on the east port of the switch for an ERPS ring named "ring1":

```
configure erps ring1 add ring-ports east 5
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

configure erps subring-mode

```
configure erps ring_name subring-mode [no-virtualChannel | virtualChannel]
```

Description

Configures sub-ring mode.

Syntax Description

<i>ring-name</i>	Alphanumeric string that identifies the <i>ERPS</i> ring.
no-virtualChannel	No Virtual Channel required to complete it's control path.
virtualChannel	Virtual Channel required to complete it's control path.

Default

N/A.

Usage Guidelines

Use this command to add or delete ERPS sub-rings.

Example

The following example configures a virtual channel for the control path:

```
configure erps ring1 subring-mode virtualChannel
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on all platforms that are running ExtremeXOS.

configure erps sub-ring

```
configure {erps} ring-name [add | delete] sub-ring-name sub_ring
```

Description

Add or delete a sub-ring to the main ring.

Syntax Description

<i>ring-name</i>	Alphanumeric string that identifies the <i>ERPS</i> ring.
add	Add sub-ring.
delete	Delete sub-ring.
<i>sub_ring</i>	Alphanumeric string identifying the ERPS sub-ring.

Default

N/A.

Usage Guidelines

Use this command to add or delete ERPS sub-rings.

Example

The following example adds sub-ring “ring2” to “ring1”:

```
configure erps ring1 add sub-ring-name ring2
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

configure erps timer guard

```
configure {erps} ring-name timer guard [ default | milliseconds ]
```

Description

Configure a guard timer to control when the node should act on received R-APS (ring automatic protection switching) messages.

Syntax Description

<i>ring-name</i>	Alphanumeric string that identifies the <i>ERPS</i> ring.
default	The default value, 500 milliseconds.
<i>milliseconds</i>	The interval for the guard timer in milliseconds, with a range of 10 to 2000.

Default

The default is 500 milliseconds.

Usage Guidelines

Use this command to configure a guard timer to control when the node should act on received R-APS messages.

Example

The following command sets the guard timer to 1000 milliseconds for an ERPS ring named “ring1”:

```
configure erps ring1 timer guard 1000
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

configure erps timer hold-off

```
configure {erps} ring-name timer hold-off [ default | milliseconds ]
```

Description

Configure a hold-off timer to control when a signal fault is relayed.

Syntax Description

<i>ring-name</i>	Alphanumeric string that identifies the <i>ERPS</i> ring.
default	The default value, 0 milliseconds.
<i>milliseconds</i>	The interval for the hold-off time in milliseconds, with a range of 0 to 10000.

Default

The default is 0 milliseconds.

Usage Guidelines

Use this command to configure a hold-off timer to control when a signal fault is relayed.

Example

The following command sets the hold-off timer to 1000 milliseconds for an ERPS ring named “ring1”:

```
configure erps ring1 timer hold-off 1000
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

configure erps timer periodic

```
configure {erps} ring-name timer periodic [ default | milliseconds ]
```

Description

Configure a periodic timer to control the interval between signal failures.

Syntax Description

<i>ring-name</i>	Alphanumeric string that identifies the <i>ERPS</i> ring.
default	The default value, 5000 milliseconds.
<i>milliseconds</i>	The interval for the periodic time in milliseconds, with a range of 2000 to 7000.

Default

The default is 5000 milliseconds.

Usage Guidelines

Use this command to configure a periodic timer to control the interval between signal failure.

Example

The following command sets the periodic timer to 6000 milliseconds for an ERPS ring named “ring1”:

```
configure erps ring1 timer periodic 6000
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

configure erps timer wait-to-block

```
configure {erps} ring-name timer wait-to-block [ default | milliseconds]
```

Description

Configure a wait-to-block timer for revertive operations on RPL owner initiated reversion.

Syntax Description

<i>ring-name</i>	Alphanumeric string that identifies the <i>ERPS</i> ring.
default	The default value, 5500 milliseconds.
<i>milliseconds</i>	The time interval to wait before restoring, with a range of 5000 to 7000 milliseconds.

Default

The default is 5500 milliseconds.

Usage Guidelines

Use this command to configure a wait-to-block timer for revertive operations on RPL owner-initiated reversion.

Example

The following command sets the wait-to-block timer to 6000 milliseconds for an ERPS ring named "ring1":

```
configure erps ring1 timer wait-to-block 6000
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

configure erps timer wait-to-restore

```
configure {erps} ring-name timer wait-to-restore [ default |  
milliseconds ]
```

Description

Configure a time interval to wait before restoring.

Syntax Description

<i>ring-name</i>	Alphanumeric string that identifies the <i>ERPS</i> ring.
default	The default value, 300000 milliseconds.
<i>milliseconds</i>	The time interval to wait before restoring, with a range of 0 to 720000 milliseconds.

Default

The default is 300000 milliseconds.

Usage Guidelines

Use this command to configure a time interval to wait before restoring.

Example

The following command sets the wait-to-restore timer to 3000 milliseconds for an ERPS ring named "ring1":

```
configure erps ring1 timer wait-to-restore 3000
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

configure erps topology-change

```
configure erps ring-name [add | delete] topology-change ring-list
```

Description

Identify the rings to which topology change events need to be propagated.

Syntax Description

<i>ring-name</i>	Alphanumeric string that identifies the <i>ERPS</i> ring.
add	Add rings/sub-rings to topology change propagation list.
delete	Delete rings/sub-rings from topology change propagation list.
<i>ring-list</i>	List of ERPS rings/sub-rings to which topology change needs to be propagated.

Default

N/A.

Usage Guidelines

Use this command to add or delete ERPS rings/sub-rings from the topology change propagation list.

Example

Example output not yet available and will be provided in a future release.

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

configure esrp add elrp-poll ports

```
configure esrp esrpDomain add elrp-poll ports [ports | all]
```

Description

Configures the ports of an *ESRP* domain where ELRP packet transmission is requested by ESRP.

Syntax Description

<i>esrpDomain</i>	Specifies an ESRP domain name.
<i>ports</i>	Specifies list of slots and ports.
all	Specifies all ports in the ESRP domain.

Default

All ports of an ESRP domain have ELRP transmission enabled.

Usage Guidelines

This command allows you to configure the ports in your network that might experience loops, such as ports that connect to master, slave, or ESRP-aware switches, to receive ELRP packets. You do not need to send ELRP packets to host ports.

Example

The following command enables ELRP packet transmission for slot 2, ports 3-5 on ESRP domain esrp1:

```
configure esrp esrp1 add elrp-poll ports 2:3-2:5
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure esrp add master

```
configure esrp esrpDomain add master vlan_name
```

Description

Adds a master VLAN to an ESRP domain.

Syntax Description

<i>esrpDomain</i>	Specifies the name of an ESRP domain.
<i>vlan_name</i>	Specifies the name of the master VLAN.

Default

N/A.

Usage Guidelines

You must configure one master VLAN for each ESRP domain. A master VLAN can belong to one ESRP domain only. An ESRP domain contains one master and zero or more member VLANs.

The master VLAN:

- Exchanges ESRP PDUs, hello messages, and data between a pair of ESRP-enabled switches.
- Contains the total number of active physical ports that are counted when determining the master ESRP domain. The switch with the highest number of active ports takes priority.

Master VLANs can have their own set of ports, and member VLANs can have a different set of ports. The state of the ESRP device determines whether the ports in the master and member VLANs are in the forwarding or blocking state.

Example

The following command adds VLAN purple to the ESRP domain esrp1 as the master VLAN:

```
configure esrp esrp1 add master purple
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure esrp add member

```
configure esrp esrpDomain add member vlan_name
```

Description

Adds a member VLAN to an ESRP domain.

Syntax Description

<i>esrpDomain</i>	Specifies the name of an ESRP domain.
<i>vlan_name</i>	Specifies the name of the member VLAN.

Default

N/A.

Usage Guidelines

You can configure zero or more member VLANs for each ESRP domain. An ESRP domain contains one master and zero or more member VLANs.

Master VLANs can have their own set of ports, and member VLANs can have a different set of ports. The state of the ESRP device determines whether the ports in the master and member VLANs are in the forwarding or blocking state.

Example

The following command adds VLAN green to the ESRP domain esrp1 as a member VLAN:

```
configure esrp esrp1 add member vlan green
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure esrp add track-environment

configure esrp *esrpDomain* **add track-environment failover** *priority*

Description

Configures an ESRP domain to track environmental failures.

Syntax Description

<i>esrpDomain</i>	Specifies an ESRP domain name.
<i>priority</i>	Specifies a number between 0 and 254. The default priority is 255. See the following "Usage Guidelines" section for more information.

Default

No environmental tracking.

Usage Guidelines

Environmental tracking tracks power supply temperature status.

If a failure is detected, the ESRP domain priority steps to the failover-priority value specified. By setting the failover priority to be lower than the normal priority of the domain, it causes the affected domain to go into slave mode.

The range of the priority value is 0 to 254. Setting the priority to 255 configures the switch to slave mode, and to be ineligible to become the master. The switch remains in slave mode even when the VLAN fails over from the current master.

To make effective use of this feature, the normal priority of the ESRP domain must be higher than the failover priority of this command.

Example

The following command enables environmental failure tracking, and specifies that the ESRP priority for ESRP domain esrp1 be set to 10 upon an environmental failure.

```
configure esrp esrp1 add track-environment failover 10
```


History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure esrp add track-iproute

```
configure esrp esrpDomain add track-iproute ipaddress/masklength
```

Description

Configures an *ESRP* domain to track a route entry in the system's routing table.

Syntax Description

<i>esrpDomain</i>	Specifies an ESRP domain name.
<i>ipaddress</i>	Specifies the IPv4 address of the route entry to be tracked.
<i>masklength</i>	Specifies the subnet of the route entry to be tracked.

Default

Disabled.

Usage Guidelines

The track-ip metric consists of the total number of tracked IPv4 routes that are up or functional.

An ESRP domain can track eight IPv4 routes.



Note

ESRP route tracking is not supported on IPv6 networks.

Example

The following command enables IPv4 route failure tracking for routes to the specified subnet:

```
configure esrp esrp1 add track-iproute 192.168.46.0/24
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure esrp add track-ping

```
configure esrp esrpDomain add track-ping ipaddress {frequency seconds}
{miss misses} {success successes}
```

Description

Configures an *ESRP* domain to track an external gateway using ping.

Syntax Description

<i>esrpDomain</i>	Specifies an ESRP domain name.
<i>ipaddress</i>	Specifies the IPv4 address of the external gateway.
frequency	Specifies setting the interval between ping requests.
<i>seconds</i>	Sets the value for the interval in seconds between ping requests. The range is 1 to 600 seconds. Default is 15.
miss	Specifies the number of consecutive ping fails required for tracking fail.
<i>misses</i>	Sets the number of consecutive failed pings to declare tracking has failed. Range is 1 to 256. Default is 4.
success	Specifies setting the number of consecutive ping successes required for tracking success.
<i>successes</i>	Sets the number of consecutive successful pings to declare tracking has succeeded. Range is 1 to 256. Default is 4.

Default

No ping tracking.

Ping successes required for tracking to succeed is 4 by default.

Ping fails required for tracking to fails is 4 by default.

The interval between ping requests is 15 seconds by default.

Usage Guidelines

The tracked-ping metric consists of the total number of stations that are successfully tracked using ping. ESRP uses an aggregate of tracked pings and traced routes to track an external gateway.

An ESRP domain can track eight stations.



Note

ESRP ping tracking is not supported on IPv6 networks.

To change any of the options for track-ping, you must delete track-ping on the ESRP domain (configure esrp esrpDomain **delete track-ping** ipaddress), and then configure it as desired.

To view track-ping options, use the command show esrp { {name} | {type [vpls-redundancy | standard]} } .

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:

```
configure esrp esrp1 add track-ping 10.207.29.17 frequency 10 miss 5
```

History

This command was first available in ExtremeXOS 11.0.

The **success** option was added in ExtremeXOS 22.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure esrp add track-vlan

```
configure esrp esrpDomain add track-vlan vlan_name
```

Description

Configures an ESRP domain to track port connectivity to a specified VLAN.

Syntax Description

esrpDomain	Specifies an ESRP domain name.
vlan_name	Specifies the VLAN to be tracked.

Default

Disabled.

Usage Guidelines

The track-vlan metric is derived from the total number of active physical ports on the VLAN being tracked by the ESRP domain.

If more than one VLAN shares a physical link, each VLAN counts the physical link.

The ESRP switch should have a higher priority number than its neighbors to ensure master election.

An ESRP domain can track one VLAN, and the tracked VLAN should not be a member of any other ESRP domain in the system.

Example

The following command enables ESRP domain esrp1 to track port connectivity to VLAN engineering:

```
configure esrp esrp1 add track-vlan engineering
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure esrp aware add selective-forward-ports

```
configure esrp domain aware add selective-forward-ports port_list {group  
group number}
```

Description

Enables selective forwarding by creating an aware port list and adds additional ports to the list.

Syntax Description

<i>domain</i>	Specifies an <i>ESRP</i> domain name.
<i>port_list</i>	Specifies the ports to be added to the aware port list.
<i>group number</i>	Specifies the ESRP group within the given domain name

Default

The group number defaults to '0'.

Usage Guidelines

An ESRP-aware switch floods ESRP PDUs from all ports in an ESRP-aware *VLAN*. This flooding creates unnecessary network traffic because some ports forward ESRP PDUs to switches that are not running the same ESRP groups. You can select the ports that are appropriate for forwarding ESRP PDUs by configuring selective forwarding on an ESRP-aware VLAN and thus reduce this excess traffic. Configuring selective forwarding creates a port list of only those ports that forward to the ESRP groups that are associated with an ESRP-aware VLAN. This ESRP-aware port list is then used for forwarding ESRP PDUs.

Use this command to create or add to an existing port list for the ESRP groups associated with an ESRP-aware VLAN.

Example

The following command configures esrp domain (d1) to forward ESRP PDUs on ports 5:1, 5:2, and 6:2.

```
configure esrp d1 aware add selective-forward-ports 5:1,5:2,6:2 group 0
```

History

This command was first available in Extreme XOS 12.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure esrp aware delete selective-forward-ports

```
configure esrp domain aware delete selective-forward-ports all | port_list  
    {group group number }
```

Description

Disables all or part of selective forwarding by deleting ports from the *ESRP*-aware port list.

Syntax Description

<i>domain</i>	Specifies an ESRP domain name.
all	Specifies that all of the ports are to be disabled.
<i>port_list</i>	Specifies the ports to be disabled from the ESRP-aware port list.
<i>group number</i>	Specifies the ESRP group within the given domain name

Default

The group number defaults to '0'.

Usage Guidelines

By configuring selective forwarding, you create an ESRP-aware port list of only those ports that forward to the ESRP groups that are associated with an ESRP-aware *VLAN*. That port list is used for forwarding ESRP PDUs from the selected ports only of an ESRP-aware switch.

Use this command to delete one or more or all of the ports from an ESRP-aware port list. Deleting all of the ports puts the domain back to the default state.

Example

The following command configures esrp domain (d1) to exclude ESRP PDUs on ports 5:1, 5:2, and 6:2.

```
configure esrp d1 aware delete selective-forward-ports 5:1,5:2,6:2 group 0
```

History

This command was first available in Extreme XOS 12.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure esrp delete elrp-poll ports

```
configure esrp esrpDomain delete elrp-poll ports [ports | all]
```

Description

Disables ELRP packet transmission on ports of an ESRP domain.

Syntax Description

<i>esrpDomain</i>	Specifies an ESRP domain name.
<i>ports</i>	Specifies list of slots and ports in the ESRP domain.
all	Specifies all ports in the ESRP domain.

Default

All ports of an ESRP domain have ELRP transmission enabled.

Usage Guidelines

If you have host ports on an ESRP domain, you do not need to send ELRP packets to those ports.

If you change your network configuration, and a port no longer connects to a master, slave, or ESRP-aware switch, you can disable ELRP transmission on that port.

Example

The following command disables ELRP packet transmission for slot 2, ports 3-5 on ESRP domain esrp1:

```
configure vlan esrp1 delete elrp-poll ports 2:3-2:5
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure esrp delete master

```
configure esrp esrpDomain delete master vlan_name
```

Description

Deletes the specifies master *VLAN* from the specified *ESRP* domain.

Syntax Description

<i>esrpDomain</i>	Specifies the name of an ESRP domain.
<i>vlan_name</i>	Specifies the name of the master VLAN.

Default

N/A.

Usage Guidelines

You must disable the ESRP domain before removing the master VLAN. To disable the ESRP domain, use the `disable esrp {esrpDomain}` command.

If you attempt to remove the master VLAN before disabling the ESRP domain, the switch displays an error message similar to the following:

```
ERROR: Failed to delete master vlan for domain "esrp1" ; ESRP is enabled!
```

If this happens, disable the ESRP domain and re-issue the `configure esrp delete master` command.

Example

The following command deletes the master VLAN purple from the ESRP domain esrp1:

```
configure esrp esrp1 delete master purple
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure esrp delete member

```
configure esrp esrpDomain delete member vlan_name
```

Description

Deletes a member VLAN from the specified ESRP domain.

Syntax Description

<i>esrpDomain</i>	Specifies the name of an ESRP domain.
<i>vlan_name</i>	Specifies the name of the member VLAN.

Default

N/A.

Usage Guidelines

None.

Example

The following command deletes the member VLAN green from the ESRP domain esrp1:

```
configure esrp esrp1 delete member vlan green
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure esrp delete track-environment

```
configure esrp esrpDomain delete track-environment
```

Description

Disables environmental failure tracking for an ESRP domain.

Syntax Description

<i>esrpDomain</i>	Specifies an ESRP domain name.
-------------------	--------------------------------

Default

No environmental tracking.

Usage Guidelines

None.

Example

The following command disables environmental failure tracking for ESRP domain esrp1:

```
configure esrp esrp1 delete track-environment
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure esrp delete track-iproute

```
configure esrp esrpDomain delete track-iproute ipaddress/masklength
```

Description

Disables route entry tracking for an *ESRP* domain.

Syntax Description

<i>esrpDomain</i>	Specifies an ESRP domain name.
<i>ipaddress</i>	Specifies the IPv4 address of a tracked route entry.
<i>masklength</i>	Specifies the subnet of a tracked route entry.

Default

Disabled.

Usage Guidelines

If you disable route tracking for a failed route, the ESRP domain recovers from the forced standby state.

If you disable route tracking for a route that is up and functional, there is no impact on the ESRP state.

Example

The following command disables tracking of routes to the specified subnet for ESRP domain esrp1:

```
configure esrp esrp1 delete track-iproute 192.168.46.0/24
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure esrp delete track-ping

```
configure esrp esrpDomain delete track-ping ipaddress
```

Description

Disables the tracking of an external gateway using ping.

Syntax Description

<i>esrpDomain</i>	Specifies an <i>ESRP</i> domain name.
<i>ipaddress</i>	Specifies the IPv4 address of the external gateway.

Default

No ping tracking.

Usage Guidelines

If you disable ping tracking for a failed ping, the ESRP domain recovers from the forced standby state.

If you disable route tracking for a successful ping, there is no impact on the ESRP state.

Example

The following command disables ping tracking for the external gateway at 10.207.29.17:

```
configure esrp esrp1 delete track-ping 10.207.29.17
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure esrp delete track-vlan

```
configure esrp esrpDomain delete track-vlan vlan_name
```

Description

Disables the tracking of port connectivity to a specified VLAN.

Syntax Description

<i>esrpDomain</i>	Specifies an <u>ESRP</u> domain name.
<i>vlan_name</i>	Specifies the VLAN to be tracked.

Default

Disabled.

Usage Guidelines

If you delete a VLAN that is down, the ESRP domain recovers from the forced standby state.

Example

The following command disables the tracking of port connectivity to VLAN engineering:

```
configure esrp esrp1 delete track-vlan engineering
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure esrp domain-id

```
configure esrp esrpDomain domain-id number
```

Description

Assigns an *ESRP* domain ID to an ESRP domain.

Syntax Description

<i>esrpDomain</i>	Specifies the name of an ESRP domain.
<i>number</i>	Specifies the number to use for the ESRP domain ID. The user-configured ID range is 4096 through 65,535.

Default

If the master *VLAN* is tagged, ESRP uses that VLANid for the ESRP domain ID. If the master VLAN is untagged, you must specify the ESRP domain ID.

Usage Guidelines

Before you enable a specific ESRP domain, it must have a domain ID. A domain ID is either a user-configured number or the VLANid of the tagged master VLAN. If you do not have a domain ID, you cannot enable ESRP on that domain.

Each switch participating in ESRP for a particular domain must have the same domain ID configured.

The number parameter range for user-configured domain IDs is 4096 through 65,535.

If the master VLAN is tagged, you can use that VLANid for the ESRP domain ID. The range for VLAN tags is 2 through 4095. Tag 1 is assigned to the default VLAN.

Example

The following command assigns the domain ID 5000 to ESRP domain esrp1:

```
configure esrp esrp1 domain-id 5000
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure esrp election-policy

```
configure esrp esrpDomain election-policy [ports > track > priority |  
ports > track > priority > mac | priority > mac | priority > ports >  
track > mac | priority > track > ports > mac | sticky > ports > track  
> priority | sticky > ports > track > priority > mac | sticky > ports
```

```
> weight > track > priority > mac | sticky > priority > mac | sticky
> priority > ports > track > mac | sticky > priority > track > ports
> mac | sticky > track > ports > priority | sticky > track > ports >
priority > mac | track > ports > priority | track > ports > priority
> mac]
```

Description

Configures the election algorithm on the switch.

Syntax Description

<i>esrpDomain</i>	Specifies an <i>ESRP</i> domain name.
ports > track > priority	Specifies that this ESRP domain should consider election factors in the following order: Active ports, tracking information, ESRP priority.
ports > track > priority > mac	Specifies that this ESRP domain should consider election factors in the following order: Active ports, tracking information, ESRP priority, MAC address. Note: This is the default election algorithm for standard mode.
priority > mac	Specifies that this ESRP domain should consider election factors in the following order: ESRP priority, MAC address.
priority > ports > track > mac	Specifies that this ESRP domain should consider election factors in the following order: ESRP priority, active ports, tracking information, MAC address.
priority > track > ports > mac	Specifies that this ESRP domain should consider election factors in the following order: ESRP priority, tracking information, active ports, MAC address.
sticky > ports > track > priority	Specifies that this ESRP domain should consider election factors in the following order: Stickiness, active ports, tracking information, ESRP priority.
sticky > ports > track > priority > mac	Specifies that this ESRP domain should consider election factors in the following order: Stickiness, active ports, tracking information, ESRP priority, MAC address.
sticky > ports > weight > track > priority > mac	Specifies that this ESRP domain should consider election factors in the following order: Stickiness, active ports, port weight, tracking information, ESRP priority, MAC address. Note: Beginning with ExtremeXOS 11.1 and later, this is the default election algorithm for extended mode.
sticky > priority > mac	Specifies that this ESRP domain should consider election factors in the following order: Stickiness, ESRP priority, MAC address.
sticky > priority > ports > track > mac	Specifies that this ESRP domain should consider election factors in the following order: Stickiness, ESRP priority, active ports, tracking information, MAC address.
sticky > priority > track > ports > mac	Specifies that this ESRP domain should consider election factors in the following order: Stickiness, ESRP priority, tracking information, active ports, MAC address.

sticky > track > ports > priority	Specifies that this ESRP domain should consider election factors in the following order: Stickiness, tracking information, active ports, ESRP priority.
sticky > track > ports > priority > mac	Specifies that this ESRP domain should consider election factors in the following order: Stickiness, tracking information, active ports, ESRP priority, MAC address.
track > ports > priority	Specifies that this ESRP domain should consider election factors in the following order: Tracking information, active ports, ESRP priority.
track > ports > priority > mac	Specifies that this ESRP domain should consider election factors in the following order: Tracking information, active ports, ESRP priority, MAC address.

Default

In extended mode, the default election algorithm is sticky > ports > weight > track > priority > mac.

In standard mode, the default election algorithm is ports > track > priority > mac.

Usage Guidelines

The election algorithm determines the order of precedence of the election factors used to determine the ESRP Master. The election factors are:

- Stickiness (sticky): the switch with the higher sticky value has higher priority. When an ESRP domain claims master, its sticky value is set to 1 (available in extended mode only).
- 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. The default priority setting is 0. A priority setting of 255 makes an ESRP switch a standby switch that remains in slave mode until you change the priority setting. We recommend this setting for system maintenance. A switch with a priority setting of 255 never becomes the master.
- MAC address (mac): the switch MAC address. A higher-number address has priority.
- Active port weight (weight)—The switch that has the highest port weight takes precedence. The bandwidth of the port automatically determines the port weight (available only in extended mode). ESRP does not count ports with a weight of 0 (known as don't count ports) regardless of ESRP running in extended or standard mode.

The election algorithm must be the same on all switches for a particular ESRP domain. The election algorithms that use sticky are and weight are available in extended mode only.

In ExtremeXOS 11.0, the extended mode default election algorithm is: sticky > ports > track > priority > mac > weight. This election algorithm is not supported in ExtremeXOS 11.1.

Factors to Consider

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 has the same ports, tracks and priority, but a higher MAC).

Any of the options with sticky can also be used to ensure that there is no failback if the original master recovers. With sticky, if an event causes the ESRP master to failover, ESRP assigns the new master with the sticky count of 1. After sticky is set on the master, regardless of changes to its neighbor's election algorithm, the new master retains its position. For example, adding active ports to the slave does not cause the new master to failback to the original master, even if the slave has more active ports than the master. Sticky algorithms provide for fewer network interruptions than non-sticky algorithms. Sticky is set on the master switch only.

ESRP re-election can occur if sticky is set on the master and a local event occurs. During this time, if the current master has lower election parameters, the backup can become the new master.

Switch Behavior

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 slave mode.

If a switch is in slave mode, it exchanges ESRP packets with other switches on that same VLAN. When a switch is in slave mode, it does not perform Layer 3 routing or Layer 2 switching services for the VLAN.

Updating the Election Algorithm

ESRP uses the default election policy for extended mode. If you have an ESRP domain operating in standard mode, the domain ignores the sticky and weight algorithms. To change the election algorithm, you must first disable the ESRP domain and then configure the new election algorithm. If you attempt to change the election algorithm without disabling the domain first, an error message appears.

To disable the ESRP domain, use the following command:

```
disable esrp {esrpDomain}
```

To modify the election algorithm, use the following command:

```
configure esrp esrpDomain election-policy [ports > track > priority |
ports > track > priority > mac | priority > mac | priority > ports >
track > mac | priority > track > ports > mac | sticky > ports > track
> priority | sticky > ports > track > priority > mac | sticky > ports
> weight > track > priority > mac | sticky > priority > mac | sticky >
priority > ports > track > mac | sticky > priority > track > ports > mac
| sticky > track > ports > priority | sticky > track > ports > priority
> mac | track > ports > priority | track > ports > priority > mac]
```

If you attempt to use an election algorithm not supported by the switch, an error message similar to the following appears:

```
ERROR: Specified election-policy is not supported!
Supported Policies:
1. sticky > ports > weight > track > priority > mac
2. ports > track > priority
3. sticky > ports > track > priority
4. ports > track > priority > mac
5. sticky > ports > track > priority > mac
6. priority > mac
```

```
7. sticky > priority > mac
8. priority > ports > track > mac
9. sticky > priority > ports > track > mac
10. priority > track > ports > mac
11. sticky > priority > track > ports > mac
12. track > ports > priority
13. sticky > track > ports > priority
14. track > ports > priority > mac
15. sticky > track > ports > priority > mac
```

Example

The following example configures the election algorithm to use tracking information as the first criteria for determining the ESRP master switch for ESRP domain esrp1:

```
configure esrp esrp1 election-policy track > ports > priority > mac
```

History

This command was first available in ExtremeXOS 11.0.

The default election algorithm for extended mode was updated to sticky > ports > weight > track > priority > mac, and the weight election factor was used in ExtremeXOS 11.1. The sticky > ports > track > priority > mac > weight election algorithm is not supported in ExtremeXOS 11.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure esrp elrp-master-poll disable

```
configure esrp esrpDomain elrp-master-poll disable
```

Description

Disables the use of ELRP by ESRP in the master state.

Syntax Description

<i>esrpDomain</i>	Specifies an ESRP domain name.
-------------------	--------------------------------

Default

Disabled.

Usage Guidelines

Use this command to disable the use of ELRP by ESRP in the master state. When you disable ELRP, the ESRP master switch no longer transmits ELRP PDUs to detect network loops.

Example

The following command disables the use of ELRP in the master state on ESRP domain elrp1:

```
configure esrp elrp1 esrp elrp-master poll disable
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure esrp elrp-master-poll enable

```
configure esrp esrpDomain elrp-master-poll enable {interval interval}
```

Description

Enables the use of ELRP by *ESRP* in the master state, and configures how often the master checks for loops in the network.

Syntax Description

<i>esrpDomain</i>	Specifies an ESRP domain name.
<i>interval</i>	Specifies how often, in seconds, successive ELRP packets are sent. The default is 1 second. The range is 1 to 64 seconds.

Default

- Use of ELRP in the master state—disabled
- Interval—1 second

Usage Guidelines

Use this command to enable the use of ELRP by ESRP in the master state. When an ESRP-enabled switch is in the master state, and you enable elrp-master-poll, the switch periodically sends ELRP PDUs at the configured interval level. If a loop is detected in the network, the transmitted PDUs are received by the switch. The ESRP master switch then transitions to the slave state to break the network loop.

We recommend that you enable both premaster and master polling when using ELRP with ESRP. To enable premaster polling, use the `configure esrp esrpDomain elrp-premaster-poll enable {count count | interval interval}`.

If you attempt to configure master polling before premaster polling, the switch displays an error message similar to the following:

ERROR: Premaster-poll should be enabled before enabling master-poll!

If this happens, first configure premaster polling followed by master polling (if required).

Specify the interval parameter to configure how often successive ELRP PDUs are sent while in the master state. If you do not specify an interval value, the default value is used.

Example

The following command enables the use of ELRP in the master state on ESRP domain elrp1:

```
configure esrp elrp1 esrp elrp-master poll enable
```

The following command configures the ESRP master to check for loops in the network every 3 seconds:

```
configure esrp elrp1 esrp elrp-master-poll enable interval 3
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure esrp elrp-premaster-poll disable

```
configure esrp esrpDomain elrp-premaster-poll disable
```

Description

Disables the use of ELRP by *ESRP* in the pre-master state.

Syntax Description

<i>esrpDomain</i>	Specifies an ESRP domain name.
-------------------	--------------------------------

Default

Disabled.

Usage Guidelines

Use this command to disable the use of ELRP by ESRP in the pre-master state. When you disable ELRP in the pre-master state, the ESRP pre-master switch no longer transmits ELRP PDUs to detect network loops prior to changing to the master state.

Example

The following command disables the use of ELRP in the pre-master state on the ESRP domain elrp1:

```
configure esrp elrp1 esrp elrp-premaster poll disable
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure esrp elrp-premaster-poll enable

```
configure esrp esrpDomain elrp-premaster-poll enable {count count |  
interval interval}
```

Description

Enables the use of ELRP by *ESRP* in the pre-master state, and configures how many times the switch sends ELRP PDUs and how often the switch sends ELRP PDUS in the pre-master state.

Syntax Description

<i>esrpDomain</i>	Specifies an ESRP domain name.
<i>count</i>	Specifies the number of times the switch sends ELRP PDUs. The default is 3. The range is 1 to 32.
<i>interval</i>	Specifies how often, in seconds, the ELRP PDUs are sent. The default is 1 second. The range is 1 to 32 seconds.

Default

- Use of ELRP in the pre-master state—disabled
- Count—3 times
- Interval—1 second

Usage Guidelines

Use this command to enable the use of ELRP by ESRP in the pre-master state to prevent network loops from occurring. When an ESRP-enabled switch is in the pre-master state (waiting to become the master), and you enable elrp-premaster-poll, the switch periodically sends ELRP PDUs at the configure level for a specified number of times. If there is a loop in the network, the transmitted PDUs are received by the switch. If this happens, the ESRP pre-master switch does not transition to the master state; rather, the switch transitions to the slave state.

We recommend that you enable both premaster and master polling when using ELRP with ESRP. To enable master polling, use the `configure esrp esrpDomain elrp-master-poll enable {interval interval}`.

If you attempt to configure master polling before premaster polling, the switch displays an error message similar to the following:

ERROR: Premaster-poll should be enabled before enabling master-poll!

If this happens, first configure premaster polling followed by master polling (if required).

If you do not specify the optional count or interval parameters, the default values are used.

If the sender does not receive packets, there is no loop in the network.

Example

The following command enables the use of ELRP—with the default settings—in the pre-master state on ESRP domain elrp1:

```
configure esrp elrp1 esrp elrp-premaster poll enable
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure esrp group

```
configure esrp esrpDomain group number
```

Description

Configures the group number to be used for the ESRP domain.

Syntax Description

<i>esrpDomain</i>	Specifies an ESRP domain name.
<i>group number</i>	Specifies the ESRP group number to which this ESRP domain should be added. The range is 0 through 31.

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 seven ESRP groups can be defined within the same networked broadcast domain. In addition, a maximum of seven distinct ESRP groups can be supported on a single ESRP switch. You can configure a maximum of 32 ESRP groups in a network.

The range for the group_number parameter is 0 through 31.

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 the same ESRP HA ports must be members of different ESRP groups.

You must first disable an ESRP domain before you modify an existing or add a new group number. If you try to modify the group number without disabling the ESRP domain, an error message similar to the following is displayed:

```
ERROR: can't change ESRP group for active domain "esrp1"!
```

To disable an ESRP domain, use the `disable esrp {esrpDomain}` command.

Example

The following command configures ESRP domain esrp1 to be a member of ESRP group 2:

```
configure esrp esrp-1 group 2
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure esrp mode

```
configure esrp mode [extended | standard]
```

Description

Configures the mode of operation for ESRP on the switch.

Syntax Description

extended	Specifies ESRP extended mode for devices running both non-Universal and modern NOS versions.
standard	Specifies ESRP standard mode for devices running modern NOS versions.

Default

The default mode is extended.

Example

The following command configures ESRP to run in standard mode:

```
configure esrp mode standard
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure esrp name

```
configure esrp esrpDomain name new-name
```

Description

Renames an existing *ESRP* domain.

Syntax Description

<i>esrpDomain</i>	Specifies the current name of an ESRP domain.
<i>new-name</i>	Specifies a new name for the ESRP domain.

Default

N/A.

Usage Guidelines

The maximum length for a name is 32 characters. Names can contain alphanumeric characters and underscores (`_`) but cannot be any reserved keywords, for example, esrp. Names must start with an alphabetical character, for example, a, Z.

You can rename an ESRP domain regardless of its current state.

Example

The following command renames ESRP domain esrp1 to esrp3:

```
configure esrp esrp1 name esrp3
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure esrp ports mode

```
configure esrp ports ports mode [host | normal]
```

Description

Configures the *ESRP* port mode for ESRP host attach.

Syntax Description

<i>ports</i>	Specifies one or more ports or slots and ports that should be configured.
host	Specifies that the ports should be configured as host ports.
normal	Specifies that the ports should be configured as normal ports.

Default

The default port mode is normal.

Usage Guidelines

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 the network to continue operation independent of ESRP status. The command sets the port to forward, allowing those ports directly attached to the slave's hosts to communicate with other hosts that are connected to the master. If you use load sharing with the ESRP HA feature, configure the load-sharing group first and then enable Host Attach on the group.

A Layer 2 connection for VLANs between ESRP switches is required.

An ESRP Host Attach port cannot be a mirroring port, software-controlled redundant port, or Netlogin port.

Example

The following command configures ports 1 through 5 on slot 3 as host ports:

```
configure esrp port 3:1-3:5 mode host
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure esrp ports no-restart

```
configure esrp ports ports no-restart
```

Description

Disables port restart for a port.

Syntax Description

<i>ports</i>	Specifies one or more ports or slots and ports.
--------------	---

Default

N/A.

Usage Guidelines

None.

Example

The following command disables port restart for ports 7-9 in slot 3 in the ESRP master domain:

```
configure esrp port 3:7-3:9 no-restart
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure esrp ports restart

```
configure esrp ports ports restart
```

Description

Configures *ESRP* to restart ports if there is a state change and the downstream switch is from another vendor.

Syntax Description

<i>ports</i>	Specifies one or more ports or slots and ports.
--------------	---

Default

N/A.

Usage Guidelines

If an ESRP domain becomes a slave, ESRP 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-enabled device. This feature allows you to use ESRP in networks that include equipment from other vendors.

If switch becomes a slave, ESRP disconnects the physical links of member ports that have port restart enabled.

An ESRP restart port cannot be a mirroring port, software-controlled redundant port, or Netlogin port.

Example

The following command enables port restart for ports 7-9 in slot 3 on the ESRP master domain:

```
configure esrp port 3:7-3:9 restart
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure esrp ports weight

```
configure esrp ports ports weight [auto | port-weight]
```

Description

Assigns the port weight for the specified *ESRP* port(s).

Syntax Description

<i>ports</i>	Specifies one or more ports or slots and ports.
auto	Specifies the switch to calculate the weight of a port based on the port's bandwidth and link speed.
<i>port-weight</i>	Specifies an ESRP port weight of 0. With a port weight of 0, the ports are not counted.

Default

The switch automatically calculates the weight of a port based on the bandwidth of the port.

Usage Guidelines

Use this command to override the automatically calculated port weight.

The port-weight parameter specifies a weight of 0. With this configuration, ESRP does not count host ports and normal ports as active. With a weight of 0, ESRP experiences fewer state changes due to frequent client activities like rebooting and unplugging laptops. A don't-count port cannot be a mirroring, software-controlled redundant port, or a Netlogin port.

For load shared ports, configure one master port in the load-share group with the port weight. A single command specifies the weight for the entire load shared group. You can specify any port from the load share group in the command. A load-shared port has an aggregate weight of all of its member ports. If you add or delete a member port (or trunk), the weight of the master load-shared port is updated. For more information about load sharing, see [Configuring Slots and Ports on a Switch](#).

Example

The following command configures port 1 on slot 3 with a weight of 0:

```
configure esrp port 3:1 weight 0
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure esrp priority

```
configure esrp esrpDomain priority number
```

Description

Configures the *ESRP* priority.

Syntax Description

<i>esrpDomain</i>	Specifies an ESRP domain number.
<i>number</i>	Specifies a number between 0 and 255.

Default

The default ESRP priority is 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 acts as master for a particular ESRP domain.

Setting the priority to 255 configures the switch to slave mode, and to be ineligible to become the master. The switch remains in slave mode even when the ESRP domain 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 ESRP domain esrp1:

```
configure esrp esrp1 priority 254
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure esrp timer hello

```
configure esrp esrpDomain timer hello seconds
```

Description

Configures the *ESRP* hello timer value.

Syntax Description

<i>esrpDomain</i>	Specifies an ESRP domain name.
<i>seconds</i>	Specifies the number of seconds between keep-alive packets. The range is 1 to 255 seconds.

Default

The default hello timer is 2 seconds.

Usage Guidelines

The timer specifies the interval, in seconds, for exchanging keep-alive packets between the ESRP switches for this ESRP domain. 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 ESRP domain across all participating switches. To see the hello settings, use the `show esrp { {name} | {type [vpls-redundancy | standard]} }` command.

The seconds range is 1 to 255.

If your configuration contains more than 2,000 ESRP VLANs and 256,000 *FDB* entries, we recommend a timer setting greater than 3 seconds.

To view the hello timer settings, use the `show esrp { {name} | {type [vpls-redundancy | standard]} }` command.

In a large ESRP configuration, the slave ESRP domain might inadvertently become the master ESRP domain. This can occur when FDB entries are flushed during a master-slave transition. To avoid this we recommend the general neighbor and hello timeout guidelines listed in [Table 7](#) on page 505, which is described in the description for the `configure esrp timer neighbor` command.

Example

The following command configures the ESRP hello timer to 4 seconds for the ESRP domain esrp1:

```
configure esrp esrp1 timer hello 4
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure esrp timer neighbor

```
configure esrp esrpDomain timer neighbor seconds
```

Description

Configures the *ESRP* neighbor timeout value.

Syntax Description

<i>esrpDomain</i>	Specifies an ESRP domain name.
<i>seconds</i>	Specifies the number of seconds after which an ESRP neighbor times out. The range is 6 to 1024 seconds.

Default

The default neighbor timeout is 8 seconds (four times the hello timer).

Usage Guidelines

The neighbor timeout specifies the amount of time that ESRP waits before considering the neighbor down. The neighbor value must be at least 3 times the hello timer value. Entering a value outside of that range generates an error message similar to the following:

```
operation Failed. Valid timer relationship "neighbor timeout >=
3*hello ; neutral timeout >= 2*hello ; premaster timeout >= 3*hello"!
```

The seconds range is 3*hello to 1024 seconds.

To view the neighbor timer settings, use the `show esrp { {name} | {type [vpls-
redundancy | standard]} }` command.

In a large ESRP configuration, the slave ESRP domain might inadvertently become the master ESRP domain. This can occur when *FDB* entries are flushed during a master-slave transition. To avoid this we recommend the general neighbor and hello timeout guidelines listed in following table.

Table 7: General Neighbor and Hello Timeout

Number of Domains	Number of VLANs	Suggested Neighbor and Hello Timeout
64 or less	1000	Use the default timer values
64	1000 to 3000	hello >=3, neighbor >=9
128	3000	hello >=4, neighbor >=12

Example

The following command configures the ESRP neighbor timeout to 14 seconds for the ESRP domain esrp1:

```
configure esrp esrp1 timer neighbor 14
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure esrp timer neutral

Configures the *ESRP* neutral timeout value.

```
configure esrp esrpDomain timer neutral seconds
```

Description

Configures the ESRP neutral timeout value.

Syntax Description

<i>esrpDomain</i>	Specifies an ESRP domain name.
<i>seconds</i>	Specifies the number of seconds after which an ESRP domain. The range is 4 to 1024 seconds.

Default

The default neutral timeout is 4 seconds (two times the hello timer).

Usage Guidelines

After you create, configure, and enable the ESRP domain, it enters the neutral state. The neutral timeout specifies the amount of time the ESRP domain stays in this temporary state before entering the slave state. The neutral value must be at least 2 times the hello timer value. Entering a value outside of that range generates an error message similar to the following:

```
operation Failed. Valid timer relationship "neighbor timeout >=
3*hello ; neutral timeout >= 2*hello ; premaster timeout >= 3*hello"!
```

The seconds range is 2*hello to 1024.

To view the neutral timer settings, use the `show esrp { {name} | {type [vpls-redundancy | standard]} }` command.

Example

The following command configures the ESRP neutral timeout to 8 seconds for the ESRP domain esrp1:

```
configure esrp esrp1 timer neutral 8
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure esrp timer premaster

```
configure esrp esrpDomain timer premaster seconds
```

Description

Configures the ESRP pre-master timeout value.

Syntax Description

<i>esrpDomain</i>	Specifies an ESRP domain name.
<i>seconds</i>	Specifies the maximum length of time, in seconds, that the transitioning master <u>VLAN</u> remains in the pre-master state. The range is 6 to 1024.

Default

The default timeout is 6 seconds (three times the hello timer).

Usage Guidelines

The premaster timer specifies how long the ESRP domain stays in the pre-master state. The pre-master timer expires if the neighbor agrees to be the slave. The premaster value must be at least three times the hello timer value. Entering a value outside of that range generates an error message similar to the following:

```
operation Failed. Valid timer relationship "neighbor timeout >=
3*hello ; neutral timeout >= 2*hello ; premaster timeout >= 3*hello"!
```

The seconds range is 3*hello-1024.

To view the pre-master timer settings, use the `show esrp { {name} | {type [vpls-redundancy | standard]} }` command.



Caution

Configure the pre-master state timeout only with guidance from Extreme Networks personnel. Misconfiguration can severely degrade the performance of ESRP and your switch.

Example

The following command configures the pre-master timeout to 10 seconds for the ESRP domain esrp1:

```
configure esrp esrp-1 timer premaster 10
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure esrp timer restart

```
configure esrp esrpDomain timer restart seconds
```

Description

Configures the *ESRP* restart timer value.

Syntax Description

<i>esrpDomain</i>	Specifies an ESRP domain name.
<i>seconds</i>	Specifies the maximum length of time, in seconds, that the neighbor ESRP switch remains in its current state during an hitless failover. The range is 2 to 1024.

Default

The default restart timer value is 2 seconds.

Usage Guidelines

The restart timer specifies the amount of time that the neighbor ESRP switch remains in its current state during a hitless failover. This timer prevent the slave ESRP switch from trying to become master during a hitless failover.

The seconds range is 2-1024.

To view the restart settings, use the `show esrp { {name} | {type [vpls-redundancy | standard]} }` command.

Example

The following command configures the restart timer value to 40 seconds for the ESRP domain esrp1:

```
configure esrp esrp-1 timer restart 40
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure failsafe-account

```
configure failsafe-account {[deny | permit] [all | control | serial |  
ssh {vr vr-name} | telnet {vr vr-name}]}
```

Description

Configures a name and password for the failsafe account, or restricts access to specified connection types.

Syntax Description

deny	Prohibits failsafe account usage over the specified connection type(s).
permit	Allows a failsafe account to be used over the specified connection type(s).
all	Specifies all connection types.
control	Specifies internal access between nodes in a SummitStack.
serial	Specifies access over the switch console port.
ssh	Specifies access using SSH on specified or all virtual routers.
telnet	Specifies access using Telnet on specified or all virtual routers.

Default

The failsafe account is always configured.

The default connection types over which failsafe account access is permitted are the same as if **permit all** is configured.

Usage Guidelines

The failsafe account is the account of last resort to access your switch.

If you use the command with no parameters, you are prompted for the failsafe account name and prompted twice to specify the password for the account. The password does not appear on the display at any time. You are not required to know the current failsafe account and password in order to change it.

If you use the command with the `permit` or `deny` parameter, the permitted connection types are altered as specified.

The failsafe account or permitted connection types are immediately saved to NVRAM on active nodes in a SummitStack.



Note

The information that you use to configure the failsafe account cannot be recovered by Extreme Networks. Technical support cannot retrieve passwords or account names for this account. Protect this information carefully.

Once you enter the failsafe account name, you are prompted to enter the password. Once you successfully log in to the failsafe account, you are logged in to an admin-level account.

Example

The following example restricts usage of the failsafe account to the series console port:

```
# configure failsafe-account deny all
# configure failsafe-account permit serial
# configure failsafe-account permit control
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure fabric attach management-vlan

```
configure fabric attach management-vlan [vlan_id | vlan_name | untagged
| none | forward [on | off] ]
```

Description

Specifies the VLAN advertised to Fabric Attach clients for them to use as the management VLAN.

Syntax Description

management-vlan	Specifies setting the VLAN advertised to Fabric Attach clients for them to use as the management VLAN. (Default is none.)
<i>vlan_id</i>	Specifies the Management VLAN ID tag (1 and 4,094).
<i>vlan_name</i>	Specifies the Management VLAN name. The VLAN must be a tagged VLAN.
untagged	Management traffic should be sent untagged.
none	No Fabric Attach Management VLAN is in use.
forward	Configures whether Fabric Attach proxy switches send Management VLAN data to clients.
on	Fabric Attach proxy switches send Management VLAN data to clients (default).
off	Fabric Attach proxy switches do not send Management VLAN data to clients.

Default

Unless configured, there is no Management VLAN by default.

If not specified, the Fabric Attach proxy switches send Management VLAN data to clients.

Usage Guidelines

Configuring the Fabric Attach Management VLAN is only relevant when operating as a Fabric Attach Server. It has no effect when operating as a client or proxy.

Fabric Attach Management VLAN forwarding configuration is relevant for all FA modes.

The management VLAN is advertised to Fabric Attach proxies and clients.

The specified VLAN configuration on the Fabric Attach server is restricted to only tagged VLANs.

Example

The following example sets the Management VLAN to a VLAN named "VLAN1" and specifies tagged traffic:

```
# configure fabric attach management-vlan VLAN1
```

History

This command was first available in ExtremeXOS 22.5.

The option to configure whether Fabric Attach proxy switches send Management VLAN data to clients was added in ExtremeXOS 30.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure fabric attach management-vlan ports

```
configure fabric attach management-vlan ports [port_list | all] forward
[on | off]
```

Description

Configures the Fabric Attach management VLAN propagation on a specific port.

Syntax Description

management-vlan	Specifies setting the VLAN advertised to Fabric Attach clients for them to use as the management VLAN. (Default is none.)
ports	Specifies the port to configure.
<i>port_list</i>	Specifies a list of ports to configure.
all	Specifies configuring all ports in the system.
forward	Configures whether Fabric Attach proxy switches send Management VLAN data to clients.
on	Fabric Attach proxy switches send Management VLAN data to clients (default).
off	Fabric Attach proxy switches do not send Management VLAN data to clients.

Default

All ports are configured to propagate the management VLAN.

When disabled, the port will not send out the management VLAN information in the element TLV Support.

Usage Guidelines

This command is only used when operating as a Fabric Attach server. It has no effect when operating as a client or proxy.

Example

The following example disables Fabric Attach port 25 from sending Management VLAN data to clients.

```
# configure fabric attach management-vlan port 25 forward off
```

History

This command was first available in ExtremeXOS 31.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure fabric attach ports

```
configure fabric attach ports [port_list | all] enable | disable
```

Description

Configures the Fabric Attach state per port.

Syntax Description

ports	Specifies the port to configure.
<i>port_list</i>	Specifies a list of ports to configure.
all	Specifies configuring all ports in the system.
enable	Specifies enabling Fabric Attach settings (default).
disable	Specifies disabling Fabric Attach settings.

Default

All ports are configured to transmit and receive Fabric Attach attributed by default.

Usage Guidelines

Use this command to enable or disable a port from adding the Fabric Attach TLVs to the LLDP packet.

Example

The following example disables Fabric Attach port 25:

```
# configure fabric attach port 25 disable
```

History

This command was first available in ExtremeXOS 31.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure fabric attach port authentication

```
configure fabric attach ports [port_list | all] authentication [ disable  
| enable | key {key | default | encrypted encrypted_key}]
```

Description

Configures Fabric Attach authentication.

Syntax Description

ports	Specify configuring ports.
<i>port_list</i>	Specifies list of ports to configure.
all	Configures all ports in the system.
authentication	Configures Fabric Attach authentication.
disable	Disable authentication setting (default).
enable	Enable authentication setting.
key	Configures Fabric Attach authentication key.
<i>key</i>	Specifies the authentication key.
default	Configures Fabric Attach authentication key to the default key. (Default when no 'key' is specified.)
encrypted	Configures Fabric Attach authentication key with encrypted key.
<i>encrypted_key</i>	Specifies the encrypted authentication key.

Default

By default, all ports are configured to authentication disabled state.

If no key is specified, the default key is used.

Usage Guidelines

When enabled, the default key is used until configured otherwise. If the authentication fails, the Fabric Attach information is dropped whether or not authentication is enabled on the receiving port.

When Fabric Attach authentication is configured on ports that are part of an , all ports on that MLAG must have the same Fabric Attach authentication configuration.

To view Fabric Attach authentication configuration, use the `show fabric attach ports [port_list | all] authentication {detail}` command. To view Fabric Attach authentication status, use the `show lldp {port [all | port_list]} neighbors {detailed}` command.

Example

The following example disables Fabric Attach authentication on all ports:

```
# configure fabric attach ports all authentication disable
```

The following example sets Fabric Attach authentication on port 1 with the default key:

```
# configure fabric attach ports 1 authentication key default
```

The following example sets Fabric Attach authentication on port 1 with the key "12345".

```
# configure fabric attach port 1 authentication key
Key: 12345
Reenter Key: 12345
```

History

This command was first available in ExtremeXOS 30.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure fabric attach uplink

```
configure fabric attach uplink [port | none]
```

Description

Configures the uplink port for and enables Fabric Attach standalone proxy operation.

Syntax Description

uplink	Uplink port for standalone proxy operation.
<i>port</i>	Enables standalone proxy operation using the specified port as the uplink.
none	Removes uplink port and disables standalone proxy operation (default).

Default

Standalone proxy mode is disabled (**none**) by default.

Usage Guidelines

Fabric Attach standalone proxy allows for Fabric Attach proxy functionality in environments without a Fabric Attach server.

The Fabric Attach standalone proxy does not send provisioning requests upstream. A Fabric Attach standalone proxy automatically accepts requests from Fabric Attach clients and assumes that the upstream network has been provisioned appropriately. Disabling Fabric Attach standalone proxy mode resets configured NSI/VLAN binding data to its default state and enables full Fabric Attach Proxy operation. In Fabric Attach standalone proxy mode, you must provide the Fabric Attach server uplink information, which is typically gathered through Fabric Attach server discovery. After you provide this information, Fabric Attach standalone proxy mode operates as if a Fabric Attach server has been discovered and is accepting NSI/VLAN binding requests. The binding clean-up is similar to a Fabric Attach server timeout event, and occurs when the static uplink is deleted and when Fabric Attach standalone proxy operation is disabled.

To confirm standalone proxy mode, use the `show fabric attach statistics` command with either the **agent** or **elements** option.

Example

The following example enables proxy mode and specifies port 10 as the uplink port:

```
# configure fabric attach uplink 10
```

The following example disables proxy mode:

```
# configure fabric attach uplink none
```

History

This command was first available in ExtremeXOS 30.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure fabric attach zero-touch-client

```
configure fabric attach zero-touch-client client [vlan [vlan_name |  
  vlan_id] [nsi nsi | isid isid] {priority [priority | dot1p]} {enable  
  | disable} | none] | enable | disable]
```

Description

Configures the Fabric Attach Zero Touch Client.

Syntax Description

zero-touch-client	Specifies Zero Touch Client.
<i>client</i>	Specifies the type of Fabric Attach client.
vlan	Specifies the VLAN on which to configure Zero Touch Clients.
<i>vlan_name</i>	Specifies the VLAN name. The VLAN must be a tagged VLAN.
<i>vlan_id</i>	Specifies the VLAN ID tag (1 and 4,094).
nsi	Network Service Identifier.
<i>nsi</i>	Specifies the Network Service Identifier. Range is 1-16777215.
insi	Individual Network Service Identifier
<i>insi</i>	Specifies the Individual Network Service Identifier. Range is 1-16777215.
priority	Specifies the Zero Touch port priority for untagged packets.
<i>priority</i>	Specifies the priority number to be used for untagged packets. Range is 0-7.

dot1p	Specifies the use of dot1p port priority for untagged packets (default).
none	Specifies no Zero Touch Client VLAN in use.
enable	Specifies enabling Zero Touch Clients.
disable	Specifies disabling Zero Touch Clients (default).

Default

Unless configured, all clients are configured to disabled state without a mapping.

Usage Guidelines

This command is only used when operating as a Fabric Attach server. It has no effect when operating as a client or proxy.

The specified VLAN configuration on the Fabric Attach server is restricted to only tagged VLANs.

Example

The following example configures the Fabric Attach Zero Touch Client "switch" with a VLAN named "V2000" and Individual Service Identifier of 2000 with a priority of 1 for untagged packets:

```
# configure fabric attach zero-touch-client switch vlan v2000 nsi 2000 priority 1
```

History

This command was first available in ExtremeXOS 31.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure fdb agingtime

```
configure fdb agingtime seconds
```

Description

Configures the *FDB* aging time for dynamic entries.

Syntax Description

agingtime	If agingtime is set to 0, all aging entries in the database are defined as static, nonaging entries.
<i>seconds</i>	Specifies the FDB aging time, in seconds. A value of 0 indicates that the entry should never be aged out. All other platforms support the value 0 (no aging) and a range of 15 to 1,000,000 seconds.

Default

300.

Usage Guidelines

If the aging time is set to 0 (zero), all dynamic entries in the database become static, nonaging entries. This means that they do not age out, but non-permanent static entries can be deleted if the switch is reset.

The software flushes the FDB table once the aging timeout parameter is reached, even if the switch is running traffic and populating addresses in the FDB table.

For ExtremeSwitching X460-G2 switches, the hardware flushes the FDB table at periods based on the configured software aging time. The actual hardware aging time does not exactly match the software aging time and can be as high as twice the configured software aging time.

Example

The following example sets the FDB aging time to 3,000 seconds:

```
# configure fdb agingtime 3000
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure fdb mac-tracking ports

```
configure fdb mac-tracking {[add|delete]} ports [port_list|all]
```

Description

Enables or disables MAC address tracking for all MAC addresses on the specified ports.

Syntax Description

add	Enables MAC address tracking for the specified ports.
delete	Disables MAC address tracking for the specified ports.
<i>port_list</i>	Specifies a list of ports on which MAC address tracking is to be enabled or disabled.
all	Specifies that MAC address tracking is to be enabled or disabled on all ports.

Default

No ports are enabled for MAC address tracking.

Usage Guidelines

MAC address tracking events on enabled ports generate EMS messages and can optionally generate *SNMP* traps.



Note
When a MAC address is configured in the tracking table, but detected on a MAC tracking enabled port, the per MAC address statistical counters are not updated.

Example

The following example enables MAC address tracking for all MAC addresses on port 2:1:

```
configure fdb mac-tracking add ports 2:1
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure fdb static-mac-move packets

```
configure fdb static-mac-move packets count
```

Description

Configures the number of EMS and *SNMP* reports that can be generated each second for MAC addresses that are duplicates of statically configured MAC addresses.

Syntax Description

<i>count</i>	Specifies the number of duplicate MAC address events that are reported each second. The range is 1 to 25.
--------------	---

Default

2.

Usage Guidelines

None.

Example

The following example configures the switch to report up to five duplicate MAC address events per second:

```
# configure fdb static-mac-move packets 5
```

History

This command was first available in ExtremeXOS 12.7.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure fdb vlan vxlan

```
configure fdb { mac_addr | broadcast | unknown-unicast | unknown-multicast } vlan vlan_name [ add | delete ] vxlan { vr vr_name } {ipaddress} remote_ipaddress
```

Description

This command allows you to add or remove remote VTEPs to a MAC address.

Syntax Description

<i>mac_addr</i>	Forwarding destination(s) for this MAC.
broadcast	Forwarding destination(s) for broadcast traffic.
unknown-unicast	Forwarding destination(s) for unknown unicast traffic.
unknown-multicast	Forwarding destination(s) for unknown multicast traffic.
add	Add to configuration.
delete	Delete from configuration.
vr	VR/VRF instance the IPv4 address is configured on.
<i>vr_name</i>	An existing VR/VRF name.
ipaddress	Configure the IP address of the remote tunnel endpoint to which the MAC needs to be bound.
<i>remote_ipaddress</i>	IPv4 address of the remote tunnel endpoint.

Default

VR-Default.

Usage Guidelines

You must first use the `[create | delete] fdb` command to add the first remote VTEP, and then issue this command to add additional remote VTEPs for the same MAC. You cannot add a remote VTEP to a static entry that has ports or blackhole configured. When the last VTEP is deleted, ExtremeXOS deletes the *FDB* entry for that MAC.

Example

```
# configure fdb 01:00:5e:00:00:01 vlan vlan101 add vxlan ipaddress 30.30.30.1
# configure fdb broadcast vlan vlan101 add vxlan vr VR-Default ipaddress 30.30.30.1
# configure fdb unknown-unicast vlan vlan101 delete vxlan ipaddress 20.20.20.1
```

History

This command was first available in ExtremeXOS 21.1.

Platform Availability

This command is supported on the ExtremeSwitching 5320, 5420, 5520, 5720 series switches, and stacks with 5320, 5420, 5520, 5720 slots only.

configure flow-redirect add nexthop

```
configure flow-redirect flow_redirect_name add nexthop ipaddress
priority number
```

Description

Adds a nexthop for the named flow redirection policy.

Syntax Description

<i>flow_redirect_name</i>	Specifies the name of the flow redirection policy.
<i>ipaddress</i>	Specifies the IPv4 or IPv6 address of a new nexthop.
<i>number</i>	Specifies the priority value for the nexthop.

Default

N/A.

Usage Guidelines

Use this command to add a new nexthop for the named flow redirection policy. You can specify an IPv4 address or an IPv6 unicast IP address (IPv6 multicast addresses are not supported). After you enter an IP address, the redirection policy only accepts addresses from the same family as the first address specified. For example, if the first IP address added is an IPv6 unicast address, you cannot add an IPv4 address to the policy.

The priority value can range from a low of 1 to a high of 4096. The nexthop with the highest priority among multiple ones is preferred as the working nexthop. When each added nexthop has the same priority, the first one configured is preferred.

Example

The following example adds a nexthop 10.1.1.1 for the flow redirection policy flow10 with a priority of 100:

```
configure flow-redirect flow10 add nexthop 10.1.1.1 priority 100.
```

History

This command was first available in ExtremeXOS 12.1.

Support for IPv6 flow-redirection policies was added in ExtremeXOS 12.7.

The maximum number of flow redirects was increased to 4096 in ExtremeXOS 16.1.

Platform Availability

This command is available for IPv4 and IPv6 flow-redirection policies on the platforms listed for the Policy Based Routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure flow-redirect delete nexthop

```
configure flow-redirect flow_redirect_name delete nexthop {ipaddress |  
all }
```

Description

Deletes a single or all nexthops for the named flow redirection policy.

Syntax Description

<i>flow_redirect_name</i>	Specifies the name of the flow redirection policy.
<i>ipaddress</i>	Specifies the IPv4 or IPv6 address of the nexthop.
all	Specifies that all configured nexthps are to be deleted.

Default

N/A.

Usage Guidelines

Use this command to delete a nexthop for the named flow redirection policy. If the deleted nexthop is the working nexthop for the policy-based routing entry, another is selected from the remaining active next hops, based on priority.

Example

The following command deletes the nexthop 10.1.1.1 from the flow redirection policy flow10:

```
configure flow-redirect flow10 delete nexthop 10.1.1.1
```

The following command deletes all configured nexthop's from the flow redirection policy exflow:

```
configure flow-redirect exflow delete nexthop all
```

History

This command was first available in ExtremeXOS 12.1.

Support for IPv6 flow-redirection policies was added in ExtremeXOS 12.7.

Platform Availability

This command is available for IPv4 and IPv6 flow-redirection policies on the platforms listed for the Policy Based Routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure flow-redirect health-check

```
configure flow-redirect flow_redirect_name health-check [ping | arp |  
neighbor-discovery]
```

Description

Configures health checking for a specific flow redirection policy.

Syntax Description

<i>flow_redirect_name</i>	Specifies the name of the flow redirection policy.
ping	Specifies ping health checking.
arp	Specifies ARP health checking for IPv4.
neighbor-discovery	Specifies Neighbor Discovery health checking for IPv6.

Default

Ping is the default.

Usage Guidelines

Use this command to configure health checking for a specific named flow redirection policy.

Example

The following command specifies arp health checking for the flow redirection policy flow10:

```
# configure flow-redirect flow10 health-check arp
```

History

This command was first available in ExtremeXOS 12.1.

Support for IPv6 flow-redirection policies was added in ExtremeXOS 12.7.

Platform Availability

This command is available for IPv4 and IPv6 flow-redirection policies on the platforms listed for the Policy Based Routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure flow-redirect nexthop

```
configure flow-redirect flow_redirect_name nexthop ip_address ping
health-check interval seconds miss number {success successes}
```

Description

Configures the ping interval, miss count, and success for a nexthop in the flow redirection policy.

Syntax Description

<i>flow_redirect_name</i>	Specifies the name of the flow redirection policy.
<i>ip_address</i>	Specifies the IPv4 or IPv6 address of the nexthop.
<i>seconds</i>	Specifies the number of seconds between pings. The default is "2".
<i>number</i>	Specifies the number of misses allowed. The default is "2".
success	Specifies a number of consecutive ping successes required to declare that a nexthop is up.
<i>successes</i>	Sets the value for the number of consecutive successful pings to declare that a nexthop is up. Range is 1 to 256. The default is 4.

Default

The default for ping interval is 2 seconds.

The default for number of misses is 2.

The default for number of successes is 4.

Usage Guidelines

Use this command to set a ping interval, miss count, and ping success. When the ping response is not received within the interval seconds * (number +1), the nexthop is considered to be dead and a new candidate is selected from the remaining active nexthops.

Example

The following command configures a ping interval of 3 seconds, miss count of 3, and success count of 3 for the nexthop 10.1.1.1 in the flow redirection policy flow 3:

```
# configure flow-redirect flow3 nexthop 10.1.1.1 ping health-check interval 3 miss 3
success 3
```

History

This command was first available in ExtremeXOS 12.1.

Support for IPv6 flow-redirection policies was added in ExtremeXOS 12.7.

The **success** option was added in ExtremeXOS 22.7.

Platform Availability

This command is available for IPv4 and IPv6 flow-redirection policies on the platforms listed for the Policy Based Routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure flow-redirect no-active

```
configure flow-redirect flow_redirect_name no-active [drop|forward]
```

Description

Configures packets to either follow the normal routing table or be dropped.

Syntax Description

<i>flow_redirect_name</i>	Specifies the name of the flow redirection policy.
drop	Specifies that the packets are to be dropped.
forward	Specifies that the packets are to follow the normal routing table.

Default

The default is forward.

Usage Guidelines

Use this command to set a drop or forward configuration for packets to be applied when all configured next hops become unreachable.

Example

The following command configures packets of the flow redirection policy flow3 to be dropped when all configured next hops become unreachable:

```
configure flow-redirect flow3 no-active drop
```

History

This command was first available in ExtremeXOS 12.1.

Support for IPv6 flow-redirection policies was added in ExtremeXOS 12.7.

Platform Availability

This command is available for IPv4 and IPv6 flow-redirection policies on the platforms listed for the Policy Based Routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure flow-redirect vr

```
configure flow-redirect flow_redirect_name vr vr_name
```

Description

Configures a virtual router for a flow redirection policy.

Syntax Description

<i>flow_redirect_name</i>	Specifies the name of the flow redirection policy.
<i>vr_name</i>	Specifies the name of the virtual router.

Default

The default virtual router is VR-Default.

Usage Guidelines

Because ACLs do not recognize the virtual router concept, one policy-based routing can be used for multiple virtual routing entries when a VLAN-based virtual router is used for one port. This configuration of a VR into a flow-redirect makes a policy-based routing work for a specific VR.

Example

The following command configures virtual router mgmt for flow redirection policy flow3:

```
configure flow-redirect flow3 vr mgmt
```

History

This command was first available in ExtremeXOS 12.1.

Support for IPv6 flow-redirection policies was added in ExtremeXOS 12.7.

Platform Availability

This command is available for IPv4 and IPv6 flow-redirection policies on the platforms listed for the Policy Based Routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure flowmon collector

```
configure flowmon collector collector_name
```

Description

Configures a created collector where Flow Monitor sends information.

Syntax Description

collector	Specifies to send flow information to a collector.
<i>collector_name</i>	Specifies the name of the created collector. Range is 32 characters.

Default

N/A.

Usage Guidelines

You must configure *src-ipv4-address*, which is a dotted decimal representation of the local address. If not configured, the collector cannot be activated.

You must configure *dst-ipv4-address*, which is a dotted decimal representation of the IPv4 address of a collector. If not configured, the collector cannot be activated.

Optional configuration parameters include the following:

- *src-udp-port* is the source port number used in the UDP header. If not specified, the switch will assign a port number.
- *dst-udp-port* is the destination port number used in the UDP header. The default value is 4,739.
- *vr-name* is the name of the virtual router used to route to the collector.
- *export-mtu* is the maximum size of a template or an IPFIX report. Default is 512 bytes, range is 512 <= <export-mtu> <= 9,216.
- *refresh-time* is the time between sends to the collector of template data in seconds (default value is 600 seconds). Minimum value is 60 seconds. This value must be coordinated with the collector by the user.

Example

The following command configures a created collector with no additional parameters:

```
# configure flowmon collector ctest dst-ipv4-address 21.1.1.100 src-ipv4-address 21.1.1.1
```

History

This command was first available in ExtremeXOS 32.2.

Platform Availability

This command is available on ExtremeSwitching 5520 and 5420 series switches.

configure flowmon group

```
configure flowmon group group_name
```

Description

Configures a created group where Flow Monitor sends information.

Syntax Description

group	Specifies the Flow Monitor group.
<i>group_name</i>	Specifies the assigned name of the Flow Monitor group. Range is 32 characters.

Default

N/A.

Usage Guidelines

Optional configuration parameters include the following:

- *max-flow-age* is the maximum age of any flow related to the group in milliseconds. The minimum value is 100ms. The default value is 60,000ms (one minute), and the maximum value is 4294967295ms (a 32-bit unsigned integer). Flow aging cannot be disabled.
- *limit* is the maximum number of flows allowed in the flow database for this group. The minimum value is 32 and the maximum value is the maximum number of flows for the switch type.
- *k-mirror* specifies the *mirror_name*. If *flow-class-id* is specified, then *k-mirror* cannot be specified. The user can clear the k-mirror using the **none** option. Traffic is mirrored to the same collector assigned to the group.
- *flow-class-id_value* is a user-assigned number between 5 and 255. If *k-mirror* is specified, then this value cannot be specified. The user can clear the value by using the **none** option.

Example

The following command configures the Flow Monitor group with the name 'src-ipv4-address':

```
# configure flowmon group src-ipv4-address
```

History

This command was first available in ExtremeXOS 32.2.

Platform Availability

This command is available on ExtremeSwitching 5520 and 5420 series switches.

configure flowmon group collector

```
configure flowmon group group_name [add | delete] collector  
                  collector_name
```

Description

Configures a relationship between a Flow Monitor collector and a group.

Syntax Description

group	Specifies the Flow Monitor group.
<i>group_name</i>	Specifies the assigned name of the Flow Monitor group. Range is 32 characters.
add	Specifies to add a collector to a group.
delete	Specifies to delete a collector from a group.
collector	Specifies to the collector.
<i>collector_name</i>	Specifies the name of the created collector. Range is 32 characters.

Default

N/A.

Usage Guidelines

Only one collector can be added to a group, but the same collector can be added to multiple groups.

The system will reject any attempt to add a collector to a group that already contains a collector.

The system will also reject any attempt to delete a collector from a group that was not previously added to the group.

Example

The following command adds a collector with the name 'src-ipv4-address' to the group with the name 'max-flow-age':

```
# configure flowmon group max-flow-age add collector src-ipv4-address
```

History

This command was first available in ExtremeXOS 32.2.

Platform Availability

This command is available on ExtremeSwitching 5520 and 5420 series switches.

configure flowmon group key

```
configure flowmon group group_name [add | delete] key key_name
```

Description

Configures a relationship between a Flow Monitor key and a group.

Syntax Description

group	Specifies the Flow Monitor group.
<i>group_name</i>	Specifies the assigned name of the Flow Monitor group. Range is 32 characters.
add	Specifies to add a key to a group.
delete	Specifies to delete a collector from a group.
key	Specifies the key.
<i>key_name</i>	Specifies the name of the created key. Range is 32 characters.

Default

N/A.

Usage Guidelines

While a key is being added to a group, the key can't be modified.

Flow Monitor creates a template key portion that matches the key. Flow Monitor then searches for a match to see if that template key portion has already been created. If not, Flow Monitor saves the new template key portion and relates it to the group. If the key has already been created, Flow Monitor checks the group to see if that template key portion is already related to the group. If a different template key portion is related to the group, then the **add** fails. If not, the template key portion (either newly created or pre-existing) is related to the key.

If the group is already enabled, the key is installed in the hardware. Keys can be added to or deleted from a group regardless of the state of the group.

When a key is to be deleted from a group, the key is disassociated from the key template portion. The key is then disassociated from the group. If the group has no more keys associated with it and the group is disabled, the group is then disassociated from the template key portion.

Example

The following command adds a key with the name 'src-ipv4-address' to the group with the name 'max-flow-age':

```
# configure flowmon group max-flow-age add key src-ipv4-address
```

History

This command was first available in ExtremeXOS 32.2.

Platform Availability

This command is available on ExtremeSwitching 5520 and 5420 series switches.

configure flowmon key ipv4

```
configure flowmon key key_name ipv4
```

Description

Configures a Flow Monitor key with IPv4 parameters.

Syntax Description

key	Specifies the Flow Monitor key.
<i>key_name</i>	Specifies the name of the key being configured with parameters. Range is 32 characters.

Default

N/A.

Usage Guidelines

Optional configuration parameters include the following:

- *src-ipv4-addr* is a source IPv4 subnetwork address given in dotted decimal notation format.
- *src-ipv4-mask* is a source IPv4 subnetwork mask value given in dotted decimal notation format.
- *dst-ipv4-addr* is a destination IPv4 subnetwork address given in dotted decimal notation format.

- *dst-ipv4-mask* is a destination IPv4 subnetwork mask value given in dotted decimal notation format.
- *port_no* is either a source or destination TCP or UDP port number.
- *protocol* is the protocol field value carried in IPv4 packets.
- *next_header* is the protocol field value carried in IPv6 packets.
- *port-list* is a list of ports given in standard CLI format, for example, 1:1-48, 3:5, and 4:6 represent ports 1 through 48 on slot 1, port 5 on slot 3, and port 6 on slot 4, respectively. Flows otherwise matching the key received on any other port will not generate a group assignment. If **all** (the default value) is specified or the *port-list* option is omitted, then the key will be installed on all available user ports that support Flow Monitor, including on any new hardware inserted into the stack.
- *key_name_other* is used with either the **before** or **after** keyword. It specifies the name of a different key from the one being configured. **before** implies that the configuring key will be higher priority than the other key, and **after** indicates that the other key will be higher priority than the one being configured. The insertion is immediately before or after the other key regardless of other keys that have been configured.

If a specific port list indicates a slot where Flow Monitor is not supported, the command will be rejected.

Example

The following command creates a key named 'k1' with no additional IPv4 parameters:

```
# configure flowmon key k1 ipv4'
```

History

This command was first available in ExtremeXOS 32.2.

Platform Availability

This command is available on ExtremeSwitching 5520 and 5420 series switches.

configure flowmon key ipv6

```
configure flowmon key key_name ipv6
```

Description

Configures a Flow Monitor key with IPv6 parameters.

Syntax Description

key	Specifies the Flow Monitor key.
<i>key_name</i>	Specifies the name of the key being configured with parameters. Range is 32 characters.

Default

N/A.

Usage Guidelines

Optional configuration parameters include the following:

- *src-ipv6-addr* is a source IPv4 subnetwork address given in dotted decimal notation format.
- *src-ipv6-mask* is a source IPv4 subnetwork mask value given in dotted decimal notation format.
- *dst-ipv6-addr* is a destination IPv4 subnetwork address given in dotted decimal notation format.
- *dst-ipv6-mask* is a destination IPv4 subnetwork mask value given in dotted decimal notation format.
- *port_no* is either a source or destination TCP or UDP port number.
- *protocol* is the protocol field value carried in IPv4 packets.
- *next_header* is the protocol field value carried in IPv6 packets.
- *port-list* is a list of ports given in standard CLI format, for example, 1:1-48, 3:5, and 4:6 represent ports 1 through 48 on slot 1, port 5 on slot 3, and port 6 on slot 4, respectively. Flows otherwise matching the key received on any other port will not generate a group assignment. If **all** (the default value) is specified or the *port-list* option is omitted, then the key will be installed on all available user ports that support Flow Monitor, including on any new hardware inserted into the stack.
- *key_name_other* is used with either the **before** or **after** keyword. It specifies the name of a different key from the one being configured. **before** implies that the configuring key will be higher priority than the other key, and **after** indicates that the other key will be higher priority than the one being configured. The insertion is immediately before or after the other key regardless of other keys that have been configured.

If a specific port list indicates a slot where Flow Monitor is not supported, the command will be rejected.

Example

The following command creates a key named 'k2' with no additional IPv6 parameters:

```
# configure flowmon key k2 ipv6
```

History

This command was first available in ExtremeXOS 32.2.

Platform Availability

This command is available on ExtremeSwitching 5520 and 5420 series switches.

configure forwarding internal-tables

```
configure forwarding internal-tables [ 12-and-13 | more [ 12 | 13-and-ipmc | routes {ipv6-mask-length [ 64 | 128] } ] ]
```

Description

Customizes the internal hardware forwarding tables based on the customer's network requirements.

Syntax Description

forwarding	Configure settings for hardware forwarding.
internal-tables	Configure settings for internal lookup tables.
12-and-13	Program the internal lookup tables for layer-2 MAC <i>FDB</i> and layer-3 hosts and IP multicast (default).
more	Configure the internal lookup tables for additional entries of specified types.
12	Program the internal lookup tables for additional layer-2 MAC FDB entries.
13-and-ipmc	Program the internal lookup tables for additional layer-3 hosts and IP multicast.
routes	Programs the internal lookup tables for additional IPv4 routes and IPv6 routes (mask 0–64) using Algorithmic Longest-Prefix Match (ALPM). This option is only available on the ExtremeSwitching 5520 series switch or stack.
ipv6-mask-length	Optimizes ALPM route capacity by choosing the maximum number of bits in the IPv6 route subnet mask length.
64	Maximizes IPv4 route capacity (default). IPv6 routes mask length: <ul style="list-style-type: none"> 0–64 bits use ALPM hardware 65–128 use ACL hardware without route sharing
128	Maximizes IPv6 route capacity for mask length 65–128 bits. All routes use ALPM hardware with route sharing.

Default

For internal tables: **12-and-13**.

For IPv6 mask length: **64**.

Usage Guidelines

Use this command to customize the internal hardware forwarding tables based on the customer's network requirements.

The ExtremeSwitching 5520 has hardware forwarding tables internal to the switch chips that can be partitioned in a flexible manner.

To display the current configuration, use the `show forwarding configuration` command.

Example

By default, the internal tables have L2 and L3 capacity whose relative size is similar to existing products. The default is:

```
# configure forwarding internal-tables l2-and-l3
```

There are three other choices. You can elect to have more L2 hardware table entries:

```
# configure forwarding internal-tables more l2
```

Or, you can choose to have more L3 unicast and multicast entries:

```
# configure forwarding internal-tables more l3-and-ipmc
```

The following example configures the switch to use ALPM to increase IPv4 and IPv6 route scaling:

```
# configure forwarding internal-tables more routes
```

The current and configured values are shown in the output of the show command:

```
# show forwarding configuration

L2 and L3 Forwarding table hash algorithm:
  Configured hash algorithm:      crc32
  Current hash algorithm:         crc32

L3 Dual-Hash configuration:
  Configured setting:             on
  Current setting:                on
  Dual-Hash Recursion Level:      1

Hash criteria for IP unicast traffic for L2 load sharing and ECMP route sharing
  Sharing criteria:               L3_L4

IP multicast:
  Group Table Compression:        on
  Local Network Forwarding:       slow-path
  Lookup-Key:                     (SourceIP, GroupIP, VlanId)

Internal lookup tables:
  Configured Setting:             more l2
  Current Setting:                l2-and-l3
NOTE: A save and reboot are required before the configured setting will take effect.

Switch Settings:
  Switching mode:                 store-and-forward

L2 Protocol:
  Fast convergence:               on

Rate Limit:
  Overhead Bytes:                 20

Fabric Flow Control:
  Fabric Flow Control:            auto
ARP and ND Settings:
  ARP Suppression Filters:        per-port
  ND Suppression Filters:         per-port
```

History

This command was first available in ExtremeXOS 15.4.

The **routes** option was added in ExtremeXOS 22.2.

The **ipv6-mask-length** option was added in ExtremeXOS 22.5.

Platform Availability

ExtremeSwitching 5520 (standalone or in a stack).

configure forwarding flow-control fabric

```
configure forwarding flow-control fabric [auto | off]
```

Description

Allows the fabric configuration to be turned off.

Syntax Description

auto	Automatically configures fabric flow control based on the priority flow control RX configuration.
off	Unconfigures the fabric flow control.

Default

Auto.

Usage Guidelines

Use this command to turn off fabric configuration or return it to the default auto mode.

Example

The following command turns off the fabric configuration:

```
configure forwarding flow-control fabric off
```

The following command returns fabric configuration to the auto mode:

```
configure forwarding flow-control fabric auto
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure forwarding hash-algorithm

```
configure forwarding hash-algorithm [crc16 | crc32]
```

Description

Modifies hardware table utilization by configuring the hash algorithm or dual-hash settings.

Syntax Description

crc16	Specifies the CRC16 hash algorithm.
crc32	Specifies the CRC32 hash algorithm. This is the default setting.

Default

In ExtremeXOS 11.5, the default hash algorithm is crc32.

In ExtremeXOS 11.4 and earlier, the default hash algorithm is crc16.

Usage Guidelines



Note

Modify the hardware table hash algorithm only with the guidance of Extreme Networks technical personnel.

The switch uses a hash algorithm to decide where to store the addresses in the hardware table. The standard, default hash algorithm works well for most systems; however, for some addresses with certain patterns, the hardware may attempt to store address information in the same section of the hardware.

If you are running ExtremeXOS 11.4 or earlier and experience a full hardware table that affects Layer 2, IP local host, and IP multicast forwarding, you see messages similar to the following in the log:

```
<Info:HAL.Ipv4Adj.Info> : adj 136.159.188.109: IP add error is Table full for new or  
newly resolved ARP, egress valid <Info:HAL.Ipv4Adj.Info> : adj 136.159.188.109: returned  
-17 for L3 table bucket 181 <Warn:HAL.Ipv4Mc.Warning> : Could not allocate a hardware  
S,G,V entry (889f4648,effffffa,70) - hardware table resource exceeded (rv=-17).
```

If you are running ExtremeXOS 11.5 or later and experience a full hardware table that affects Layer 2, IP local host, and IP multicast forwarding, you see messages similar to the following in the log:

```
<HAL.Ipv4Adj.L3TblFull> MSM-A: IPv4 unicast entry not added. Hardware L3 Table full.
```

In the previously described situations, you can configure a different hash algorithm to select a different section of the hardware to store addresses. You must save your configuration and reboot the switch to modify the hash algorithm used by the hardware table. Typically, the dual-hash feature improves hash utilization. You must save your configuration and reboot the switch to turn dual-hash on or off.

Upgrading to ExtremeXOS 11.5

When you upgrade to ExtremeXOS 11.5, the hash algorithm automatically becomes crc32. For example, if you saved a configuration using an image from ExtremeXOS 11.4 or earlier with the hash algorithm set to crc16, when ExtremeXOS 11.5 loads, the hash algorithm becomes crc32. To change the hash algorithm to crc16, use the `configure forwarding hash-algorithm crc16` and save your switch configuration.

Example

The following example modifies the hardware table hash algorithm to crc16:

```
configure forwarding hash-algorithm crc16
```

The switch displays the following message to describe the change and to prompt you to save your configuration and reboot the switch:

```
Configured hash alorithm has been changed to 'crc16' with L3 dual-hash support 'on' for
applicable HW.
```

```
Warning: This command will only take effect after a save and reboot
```

The switch displays the following message:

```
Configured hash algorithm has been changed to 'crc32' with L3 dual-hash support 'off' for
applicable HW.
```

```
Warning: This command will only take effect after a save and reboot.
```

To display the results, use the `show forwarding configuration` command.

History

This command was first available in ExtremeXOS 11.3.2.

The default hash algorithm was changed to crc32 in ExtremeXOS 11.5.

Platform Availability

This command is available only on all platforms.

configure forwarding hash-recursion-level

```
configure forwarding hash-recursion-level 0-3
```

Description

Modifies hardware table utilization by configuring the dual hashing recursion level.

Syntax Description

0-3	Sets the maximum number of L3 hash buckets to modify to make room for a new entry.
-----	--

Default

The default is “1.”

Usage Guidelines

This command allows you to select the dual hashing “recursion level” for hardware with the dual-hash feature. The setting applies only if dual-hash is configured or defaulted to “on” using the configure forwarding hash-algorithm command.

The configured recursion level is the maximum number of existing hash entries to move in an attempt to add a new hash entry. A higher recursion level may provide better hash utilization at the expense of additional CPU processing. This command does not require a system reboot. However, the new recursion level takes effect only for addresses added after the command is issued.

Example

The following command modifies the dual-hash recursion level to modify up to two L3 hash buckets in an attempt to add a new entry:

```
configure forwarding hash-recursion-level 2
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure forwarding ipmc all

```
configure forwarding ipmc all [flood | learn]
```

Description

Enables you to forward packets for the mDNS, LLMNR, and UPnP protocols to the VLAN for learning.

Syntax Description

flood	Specifies to forward all (mDNS, LLMNR, and UPnP) to VLAN. (Default).
learn	Specifies to send all (mDNS, LLMNR, and UPnP) to the CPU for learning.

Default

flood

Usage Guidelines

Use this command to enable mDNS, LLMNR, and UPnP for flood or learn state.

When **flood** to VLAN is enabled, an ACL filter is installed, and one ACL entry is consumed for each of the three protocols.

When in **learn** mode, no ACL filters are installed, and no ACL resources are consumed.

Use the **llmnr**, **mdns**, and **upnp** keywords to specify separate protocols.

Example

The following command enables all three protocols to flood to VLAN:

```
configure forwarding ipmc all flood
```

History

This command was first available in ExtremeXOS 32.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure forwarding ipmc compression

```
configure forwarding ipmc compression {group-table | off}
```

Description

Enables or disables compression of entries in the IP multicast group table to facilitate improved IP multicast scaling.

Syntax Description

group-table	Enables compression.
off	Disables compression.

Default

group-table.

Usage Guidelines

Compression of IP multicast group table entries allows the switch to process more multicast traffic using the faster switch hardware instead of the relatively slower switch software. Compression requires additional processing. Disable this feature if you suspect a problem exposed by IP multicast compression.

When you enable or disable this feature, all IP multicast entries are flushed, and this can result in a temporary loss of multicast traffic while the IP multicast entries are relearned.

To display the compression feature configuration, enter the command:

```
show forwarding configuration
```

Example

The following command disables compression:

```
configure forwarding ipmc compression off
```

History

This command was first available in ExtremeXOS 12.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure forwarding ipmc llmnr

```
configure forwarding ipmc llmnr [flood | learn]
```

Description

Enables you to forward packets for the LLMNR protocol to the VLAN for learning.

Syntax Description

flood	Specifies to forward LLMNR to VLAN. (Default).
learn	Specifies to send LLMNR to the CPU for learning.

Default

flood

Usage Guidelines

Use this command to enable LLMNR to a flood or learn state.

When **flood** to VLAN is enabled, an actual ACL filter is installed, and one ACL entry is consumed.

When in **learn** mode, no ACL filters are installed, and no ACL resources are consumed.

Example

The following command enables LLMNR to flood to VLAN:

```
configure forwarding ipmc llmnr flood
```

History

This command was first available in ExtremeXOS 32.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure forwarding ipmc local-network-range

```
configure forwarding ipmc local-network-range [fast-path | slow-path]
```

Description

Sets how forwarding of packets to local network IP multicast addresses (224.0.0.x) is handled.

Syntax Description

fast-path	Specifies fast-path forwarding. Fast-path forwarding dictates that packets traversing the switch do not require processing by the CPU. Fast path packets are forwarded entirely by ASICs and are sent at wire speed rate. This consumes additional system <i>ACL</i> , per-port or per- <i>VLAN</i> , depending on <i>configure igmp snooping filters</i> [per-port per-vlan] selections.
slow-path	Specifies slow-path forwarding (default). Packets are processed by the CPU.

Default

Slow-path forwarding is the default configuration.

Example

The following example sets up fast-path forwarding for local network IP multicast addresses:

```
configure forwarding ipmc local-network-range fast-path
```

History

This command was first available in ExtremeXOS 15.3.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure forwarding ipmc lookup-key

```
configure forwarding ipmc lookup-key [group-vlan | source-group-vlan |
mac-vlan | mixed-mode]
```

Description

Enables you to choose the lookup-key for multicast forwarding.

Syntax Description

group-vlan	Specifies that IP multicast forwarding database entries are programmed as (*,GroupIP,VlanId).
source-group-vlan	Specifies that IP multicast forwarding database entries are programmed as (SourceIP, GroupIP, VlanId). (Default).
mac-vlan	Specifies that IP multicast forwarding database entries are programmed as (Mac, VlanId).
mixed-mode	Specifies that IP multicast forwarding database entries are programmed as follows: L3 cache entries (PIM/MVR/PVLAN) use source-group-vlan; L2 cache entries (<u>IGMP</u> /MLD/PIM snooping) use mac-vlan.

Default

source-group-vlan.

Usage Guidelines

Use this command to choose the lookup-key for multicast forwarding. The following restrictions apply to this command:

The `configure forwarding ipmc lookup-key mac-vlan` command is disallowed under the following conditions.

- If IPMC forwarding is enabled on at least on one VLAN
- If MVR is enabled either globally or on a VLAN

Similarly, enabling the above two features are disallowed,when the ipmc lookup-key is **mac-vlan**.

The following warning message is displayed when the **mac-valn** option is specified:

```
Warning: Usage of multicast IP addresses that could result in overlapping MAC
addresses should be avoided. Example: Using 225.1.1.1, 226.1.1.1 and 225.129.1.1 should
be avoided. Either one of the addresses could be used. Using multicast with PVLAN
should be avoided with this forwarding option.
```

- **Mixed-mode:** configure forwarding ipmc lookup-key mixed-mode
- The configure igmp snooping forwarding-mode [**group-vlan** | **source-group-vlan**] command was introduced to support (*, G, V) forwarding before the IPMC compression feature was introduced. Because we are introduced IPv6 multicast support in ExtremeXOS 15.2, this command is deprecated, and the new configure forwarding ipmc lookup-key command now covers both IPv4 and IPv6.

The following warning message appears when the mixed mode option is specified:

Warning: Usage of multicast IP addresses that could result in overlapping MAC addresses should be avoided for snooping (IGMP/MLD/PIM snooping) controlled traffic.

Example: Using 225.1.1.1, 226.1.1.1 and 225.129.1.1 should be avoided. Either one of the addresses could be used.

Example

The following command specifies that IP multicast forwarding database entries are programmed as (*,GroupIP,VlanId):

```
configure forwarding ipmc lookup-key group-vlan
```

To display the ipmc lookup-key configuration, enter the command:

```
show forwarding configuration
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure forwarding ipmc mdns

```
configure forwarding ipmc mdns [flood | learn]
```

Description

Enables you to forward packets for the mDNS protocol to the VLAN for learning.

Syntax Description

flood	Specifies to forward mDNS to VLAN. (Default).
learn	Specifies to send mDNS to the CPU for learning.

Default

flood

Usage Guidelines

Use this command to enable mDNS for flood or learn state.

When **flood** to VLAN is enabled, an actual ACL filter is installed, and one ACL entry is consumed.

When in **learn** mode, no ACL filters are installed, and no ACL resources are consumed.

Example

The following command enables mDNS to flood to VLAN:

```
configure forwarding ipmc mdns flood
```

History

This command was first available in ExtremeXOS 32.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure forwarding ipmc upnp

```
configure forwarding ipmc upnp [flood | learn]
```

Description

Enables you to forward packets for the UPnP protocol to the VLAN for learning.

Syntax Description

flood	Specifies to forward UPnP to VLAN. (Default).
learn	Specifies to send UPnP to the CPU for learning.

Default

flood

Usage Guidelines

Use this command to enable UPnP for flood or learn state.

When **flood** to VLAN is enabled, an actual ACL filter is installed, and one ACL entry is consumed.

When in **learn** mode, no ACL filters are installed, and no ACL resources are consumed.

Example

The following command enables UPnP to flood to VLAN:

```
configure forwarding ipmc upnp flood
```

History

This command was first available in ExtremeXOS 32.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure forwarding L2-protocol fast-convergence

```
configure forwarding L2-protocol fast-convergence on | off
```

Description

Configures the switch to flooding the unicast traffic during L2 protocol convergence.

Syntax Description

on	Used to avoid flooding the unicast traffic during L2 protocol convergence. (default)
off	Used to Temporarily flooding unicast traffic during L2 protocol convergence.

Default

On.

Usage Guidelines

Use this command to influence the L2-protocol convergence when topology changes in the network to minimize the congestion.

Example

The following command will influence the L2-Protocol control traffic:

```
configure forwarding L2-protocol fast-convergence off
```

History

This command was first available in ExtremeXOS 15.1.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure forwarding rate-limit overhead-bytes

```
configure forwarding rate-limit overhead-bytes overhead_bytes
```

Description

This command allows you to select the number of overhead bytes that will be included in the rate calculation.

Syntax Description

<code>rate-limit</code>	Rate limiting features.
<code><i>overhead_bytes</i></code>	Number of overhead bytes used in rate-limit and meter calculations.

Default

20 bytes to include the preamble and inter-frame gap.

Example

The following example displays the output of the show forwarding configuration command with the rate limit information included.

```
L2 and L3 Forwarding table hash algorithm:
  Configured hash algorithm:      crc32
  Current hash algorithm:         crc32

L3 Dual-Hash configuration:
  Configured setting:             on
  Current setting:               on
  Dual-Hash Recursion Level:      1

Hash criteria for IP unicast traffic for L2 load sharing and ECMP route sharing
  Sharing criteria:              L3_L4

IP multicast:
  Group Table Compression:        on
  Local Network Forwarding:       slow-path
  Lookup-Key:                    (SourceIP, GroupIP, VlanId)

Internal lookup tables:
  Configured Setting:            l2-and-l3
  Current Setting:              l2-and-l3

Switch Settings:
  Switching mode:                store-and-forward

L2 Protocol:
  Fast convergence:              on

Rate Limit:
  Overhead Bytes:                20
```

```

Fabric Flow Control:
  Fabric Flow Control:          auto

ARP and ND Settings:
  ARP Suppression Filters:      per-port
  ND Suppression Filters:       per-port

```

History

This command was first available in ExtremeXOS 21.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure forwarding sharing

```
configure forwarding sharing [L3 | L3_L4]
```

Description

Identifies the fields that are used to select *ECMP* routes and load-sharing group ports.

Syntax Description

L3	Uses only Layer 3 IP addresses to select ECMP routes and load-sharing ports.
L3_L4	Uses Layer 3 IP addresses and Layer4 TCP/UDP port numbers, if present, to select ECMP routes and load-sharing ports.

Default

L3_L4.

Usage Guidelines

This command configures the criteria used to select ECMP routes and load-sharing group ports.

For ECMP routes, the configured criteria selects the next hop gateway. The L3 option uses only the source and destination IP addresses to select the next hop gateway. The L3_L4 option uses the Layer4 TCP or UDP port and the source and destination IP addresses to select the next hop gateway.

For load-sharing groups (link aggregation groups), the configured criteria selects the load-sharing group port. The load-sharing groups can be configured to use the following address-based algorithms:

- L2—Specifies port selection based on Layer 2 information.
- L3—Specifies port selection based on Layer 3 information.
- L3_L4—Specifies port selection based on Layer 3 and Layer4 information.

This command affects all the load-sharing groups that use either the L3 or L3_L4 link aggregation algorithm. If the L3 option is specified, all the load-sharing groups that are configured with either the L3 or the L3_L4 address-based link aggregation algorithm use just the Layer 3 IP addresses for the egress port selection. Similarly if the L3_L4 option is specified, all the load-sharing groups that are configured with either L3 or L3_L4 address-based link aggregation algorithm use the Layer 3 IP addresses and Layer4 port number for the egress port selection.

Selecting the L3 option over L3_L4 can be useful in a network where IP fragments are present, since only the first fragment contains the Layer4 TCP or UDP port number. If the L3 option is selected, all IP fragments in a given TCP or UDP session use the same ECMP gateway or load-sharing group port, potentially avoiding inefficient packet reordering by the destination. If IP fragments are not prevalent, better traffic distribution can be achieved by selecting L3_L4.

To display the forwarding sharing feature configuration, enter the command: `show forwarding configuration`

Example

The following example modifies the sharing selection criteria to use just the Layer 3 IP addresses:

```
configure forwarding sharing L3
```

The following example modified the sharing selection criteria to use the Layer 3 and Layer 4 information:

```
configure forwarding sharing L3_L4
```

History

This command was first available in ExtremeXOS 11.6.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure forwarding suppression filters

```
configure forwarding iparp suppression filters [per-port |per-vlan]
```

Description

This command controls the way the hardware filters are installed for VXLAN ARP suppression.

Syntax Description

iparp	Selects IP ARP.
suppression	ARP suppression. Requests may be proxied.
filters	Control the way ARP suppression hardware filters are installed.

per-port	Install ARP suppression hardware filters on a per-port basis (default).
per-vlan	Install ARP suppression hardware filters on a per-VLAN basis.

Default

By default, **per-port** option is assumed.

Example

The following example sets IP ARP suppression filtering per-VLAN:

```
configure forwarding iparp suppression filters per-vlan
```

History

This command was first available in ExtremeXOS 22.3.

Platform Availability

This command is supported on the ExtremeSwitching 5420, 5520 series switches, and stacks with 5320, 5420, 5520, 5720 slots only.

configure forwarding vpex ipmc replication

```
configure forwarding vpex ipmc replication [ controlling-bridge | bpe ]
```

Description

Configures Extended Edge Switching optimized IP multicast replication mode.

Syntax Description

forwarding	Designates configuring hardware capabilities.
vpex	Designates configuring Extended Edge Switching capabilities.
ipmc	Designates configuring IP multicast forwarding settings.
replication	Designates configuring the way hardware replicates IP multicast packets.
controlling-bridge	Sets IP multicast replication to be performed at the controlling bridge (CB) (default).
bpe	Sets IP multicast replication to be performed at the bridge port extender (BPE).

Default

By default, IP multicast replication is done on the CB.

Usage Guidelines

IP multicast (IPMC) replication involves generating multiple copies of incoming IPMC traffic to subscribed receivers. Without IPMC replication, in Extended Edge Switching, if there are “N” receivers on the BPEs, “N” copies are generated by the CB, which is not efficient. IPMC replication allows the CB to send one copy of the packet to the BPE, and then let the BPE do the replication.

You can configure the replication mode with this command. The default replication is at the CB, but you can configure replication at the BPE. You might want to switch to CB-based IPMC replication for debugging in the event of a problem or scalability, because replication on the BPE uses a hardware resource on the CB that is limited.

To see the current IP multicast replication mode, use the command `show forwarding configuration`.

Example

The following example configures the IPMC replication to BPE mode:

```
# configure forwarding vpex ipmc replication bpe
```

History

This command was first available in ExtremeXOS 30.6.

Platform Availability

This command is available on ExtremeSwitching 5420 and 5520 series switches.

configure forwarding vpex vlan-port-filter

```
configure forwarding vpex vlan-port-filter [hash-table | port-group]
```

Description

Selects the way VLAN membership is implemented for Extended Edge Switching extended ports.

Syntax Description

vpex	Specifies Extended Edge Switching.
vlan-port-filter	Select hardware mechanism to enforce VLAN port membership.
hash-table	Use hash table for VLAN port membership when different VLANs do not share many ports (default).
port-group	Use port group for VLAN port membership when different VLANs share many ports and there is a requirement for large VLAN scale.

Default

Hash table is the default behavior.

Usage Guidelines

In Extended Edge Switching hardware, the extended ports are represented as virtual ports. The VLAN membership of extended ports can be implemented in two ways:

- Hash table with VLAN and virtual port as key. Note that hash tables can lead to hash collisions at higher scale. (Default)
- Virtual port group. Programming the same group number in the VLAN table and virtual port table indicates membership. The hardware has 64 virtual port groups. You should select this option if many VLANs share the same extended ports.

Note that changing this configuration at run time could result in temporary loss of traffic while the tables are reprogrammed. It is preferable to identify which option works best for the particular topology and leave the setting unchanged during run time or schedule the change during a maintenance window.

To see what setting you have selected with this command, see [show forwarding configuration](#) on page 2670.

Example

The following example selects a virtual port group to define VLAN membership:

```
# configure forwarding vpex vlan-port-filter port-group
```

History

This command was first available in ExtremeXOS 22.6.

Platform Availability

This command is available on ExtremeSwitching 5420 and 5520 series switches.

configure identity-management role

```
configure identity-management role role_name { tag [ tag | none ] } {vr  
  vr_name | none ] }
```

Description

This command defines VLAN/VR membership to an identity management role.

Syntax Description

<i>role_name</i>	Name of the role.
tag	VLAN tag for dynamic VLAN creation for this role.
<i>tag</i>	VLAN tag between 1 and 4094.

vr	Virtual router name for dynamic VLAN creation for this role.
<i>vr_name</i>	Virtual router name.
none	None.

Default

N/A.

Usage Guidelines

Use this command to configure VLAN tag and the VR in which the dynamic VLAN has to be created for a role. By default the dynamic VLAN is created in *VR-Default* if the VR is not configured. The identity is placed in the base VLAN if no VLAN tag is configured for this role. The configured VLAN tag and VR can be set to none to unconfigure the same. *VR-Mgmt* is not allowed to configure. The VLAN tag and VR is applicable only to the user created roles.

Example

The following example configures role "r1" and tag 100:

```
# configure identity-management role "r1" tag 100 vr "VR-Default"
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure identity-management role-based-vlan

```
configure identity-management role-based-vlan [add | delete] ports  
[port_list | all]
```

Description

This command defines

Syntax Description

role-based vlan	Associates the identity to a specific <i>VLAN</i> based on the identity's role.
add	Adds ports to the Identity Management role-based vlan enabled portlist.
delete	Deletes ports from the Identity Management role-based vlan enabled portlist.

ports	Configures Identity Management role-based VLAN on ports.
<i>port_list</i>	Configures Identity Management role-based VLAN on specified port list.
all	Configures Identity Management role-based VLAN on all ports.

Default

N/A.

Usage Guidelines

Use this command to configure the role-based VLAN feature for Identity Management enabled ports. This command requires the ports to be part of a base VLAN. Enabling role-based VLAN on Identity Management enabled ports allows the identity to be placed in the correct VLAN mapped to the role as configured by the administrator.



Note

You cannot enable the Identity Manager role-based VLAN feature on Netlogin enabled ports.

Example

The following example configures Identity Management on ports 1-3, and 5.

```
# configure identity-management role-based-vlan add ports 1-3,5
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure identity-management access-list

```
configure identity-management access-list source-address [mac | ip]
```

Description

Configures the access-list source-address type.

Syntax Description

mac	Specifies MAC addresses.
ip	Specifies IP addresses.

Default

MAC addresses.

Usage Guidelines

The identity management feature can install ACLs for identities based on the source MAC or source IP address. By default the MAC address of the identity is used to install the ACLs. Every network entity has a MAC address, but not all network devices have an IP address, so we recommend that you use the default mac selection to install ACLs for network entities based on the source MAC address.

You must disable the identity management feature with the `disable identity-management` command before you use this command.

Example

The following command configures the identity management feature to use MAC-based ACLs:

```
* Switch.4 # configure identity-management access-list source-address mac
```

History

This command was first available in ExtremeXOS 12.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure identity-management blacklist

```
configure identity-management blacklist add [mac mac_address {macmask}
| ip ip_address {netmask} | ipNetmask] | user user_name] configure
identity-management blacklist delete [all | mac mac_address {macmask}
| ip ip_address {netmask} | ipNetmask] | user user_name]
```

Description

Adds or deletes an entry in the identity manager blacklist.

Syntax Description

add	Adds the specified identity to the blacklist.
delete	Deletes the specified identity from the blacklist.
all	Specifies that all identities are to be deleted from the blacklist. This option is available only when the delete attribute is specified.
<i>mac_address</i>	Specifies an identity by MAC address.
<i>macmask</i>	Specifies a MAC address mask. For example: FF:FF:FF:00:00:00.

<i>ip_address</i>	Specifies an identity by IP address.
<i>netmask</i>	Specifies a mask for the specified IP address.
<i>ipNetmask</i>	Specifies an IP network mask.
<i>user_name</i>	specifies an identity by user name.

Default

N/A.

Usage Guidelines

The software supports up to 512 entries in the blacklist. When you add an identity to the blacklist, the switch searches the whitelist for the same identity. If the identity is already in the whitelist, the switch displays an error.

It is possible to configure an identity in both lists by specifying different attributes in each list. For example, you can add an identity username to the blacklist and add the MAC address for that user's laptop in the whitelist. Because the blacklist has priority over the whitelist, the username is denied access to the switch from all locations.

If you add a new blacklist entry that is qualified by a MAC or IP address, the identity manager does the following:

- Reviews the identities already known to the switch. If the new blacklist entry is an identity known on the switch, all existing ACLs (based on user roles or whitelist configuration) for the identity are removed.
- When a blacklisted MAC-based identity is detected or already known, a Deny All *ACL* is programmed for the identity MAC address for the port on which the identity is detected.
- When a blacklisted IP-based identity is detected or already known, a Deny All ACL is programmed for the identity IP address for the port on which the identity is detected.
- The ACL for blacklisted MAC and IP addresses precedes any ACLs based on user names (including Kerberos snooping) that may have been previously configured on the port. This ensures that a Kerberos exchange cannot complete when initiated for blacklisted identities.

If you add a new blacklist entry that is qualified by a username (with or without a domain name), the identity manager does the following:

- Reviews the identities already known to the switch. If the new blacklist entry is an identity known on the switch, a Deny All ACL is programmed for the identity MAC address on all ports to which the identity is connected.
- When a new blacklisted username-based identity accesses the switch, a Deny All ACL is programmed for the identity MAC address on the port on which the identity was detected.
- The ACL for a blacklisted username follows any ACLs based on Kerberos snooping. This ensures that a Kerberos exchange for another user can complete when initiated from the same MAC address.



Note

Identity manager programs ingress ACLs. Blacklisted devices can receive traffic from the network, but they cannot send traffic into the network.

Deny All ACLs for blacklisted entries exist as long as the identity remains in the identity manager database.

If you delete an identity from the blacklist, identity manager checks to see if the identity is in the local database. If the identity is known to the switch, the switch does the following:

- Removes the Deny All ACL from the port to which the identity connected.
- Initiates the role determination procedure for the switch port to which the known identity connected. This ensures that the appropriate role is applied to the identity that is no longer blacklisted.

**Note**

The role determination process can trigger an LDAP refresh to collect identity attributes for role determination.

Example

The following command adds a MAC address to the blacklist:

```
* Switch.4 # configure identity-management blacklist add mac 00:01:05:00:03:18
```

The following command deletes a user name from the blacklist:

```
* Switch.5 # configure identity-management blacklist delete user bill_jacob@b.com
```

History

This command was first available in ExtremeXOS 12.7.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure identity-management database memory-size

```
configure identity-management database memory-size Kbytes
```

Description

Configures the maximum amount of memory that is allocated to the identity management database.

Syntax Description

<i>Kbytes</i>	Specifies the maximum amount of memory to be used for maintaining identity information. The range is 64 to 49152 KB.
---------------	--

Default

512 KB.

Usage Guidelines

If the current memory usage is higher than the memory size specified in the configure identity-management database memory-size command, the command is not successful and a warning message appears. The message indicates that the current memory usage level is higher than the configured level and that the memory can be freed only when existing identities log out or disconnect.

Example

The following command allocates 4096 kilobytes to the identity management database:

```
* Switch.4 # configure identity-management database memory-size 4096
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure identity-management detection

```
configure identity-management detection [on | off] [fdb | iparp |  
ipsecurity | kerberos | lldp | netlogin | all] ports [port_list |  
all]
```

Description

This command provides the administrator a way to enable/disable the detection of the identities that are triggered through any of the following protocols:

- FDB
- IPARP
- IPSecurity DHCP Snooping
- LLDP
- Netlogin
- Kerberos

Syntax Description

detection	Detection of the identities.
on	Detection of identities on.

off	Detection of identities off.
fdb	FDB identities.
iparp	IPARP identities.
ipsecurity	Identities detected through DHCP snooping entries.
kerberos	Kerberos identities.
lldp	LLDP identities.
all	All identities.

Default

On.

Usage Guidelines

The identity manager detects the identities using the following protocols:

- FDB
- IPARP
- IPSecurity DHCP Snooping
- LLDP
- Netlogin
- Kerberos

By default, Identity Management detects identities through all the above mentioned protocols.

This feature provides the administrator a way to enable/disable the detection of the identities that are triggered through any of the above said protocols. The administrator can control the identity detection through any of the protocol trigger at the port level. This configuration can be applied to identity management enabled ports only. ExtremeXOS displays an error if this configuration is applied for the identity management disabled ports.



Note

All types of Netlogin identity will not be detected if the netlogin detection is disabled. Enabling Kerberos identity detection will not create identities for the previously authenticated Kerberos clients.

Example

```
* Slot-1 Stack.1 # configure identity-management detection off fdb ports 1:3-6
* Slot-1 Stack.2 # configure identity-management detection off ipsecurity ports 1:3-6
* Slot-1 Stack.3 # configure identity-management detection off kerberos ports 1:1, 2:5-8
* Slot-1 Stack.4 # configure identity-management detection off netlogin ports 1:1-24,
2:1-24
The effect of these commands can be seen by issuing the show identity-management command
* Slot-1 Stack.5 # show identity-management
Identity Management : Enabled
Stale entry age out (effective) : 180 Seconds (180 Seconds)
```

```

Max memory size : 512 Kbytes
Enabled ports : 1:1-24, 2:1-24
FDB Detection Disabled ports : 1:3-6
IPARP Detection Disabled ports : None
IPSecurity Detection Disabled ports : 2:1
Kerberos Detection Disabled ports : 1:1, 2:5-8
LLDP Detection Disabled ports : None
Netlogin Detection Disabled ports : 1:1-24, 2:1-24
SNMP trap notification : Enabled
Access list source address type : IP
Kerberos aging time (DD:HH:MM) : 00:08:00
Kerberos force aging time (DD:HH:MM) : None
Valid Kerberos servers : none configured(all valid)

```

History

This command was first available in ExtremeXOS 15.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure identity-management greylist

```

configure identity-management greylist add user username identity-
management greylist delete [all | user username]

```

Description

This command enables a network administrator to choose usernames whose identity is not required to be maintained. These user names are added to greylist. Identity Management module does not create an identity when greylist users log in.

Syntax Description

<i>username</i>	Specifies an identity by user name.
-----------------	-------------------------------------

Default

N/A.

Usage Guidelines

The software supports up to 512 entries in greylist. Administrator can configure username as part of greylist. When such configuration takes place, identity manager takes following action.

- Checks if the same entry is present in blacklist/whitelist. If yes, command is rejected with appropriate error message.
- Checks if this entry is ineffective because of existing entries in blacklist/whitelist. During this check, precedence of greylist is also taken into account.
 - E.g: New entry being configured into greylist is: Richard@corp. Assume blacklist has higher precedence and it has an entry "Richard". In this case, new entry is ineffective and the configuration is rejected giving the details.
- If no conflict is found, greylist is updated.
- IDM checks if any existing identity matches the new entry in greylist. If match is found, location/identity will be deleted and unknown identity is created with the same MAC.

If greylist user is the only user logged into the device, unknown identity is created and user is kept in unauthenticated role. However if actual user is present along with greylist user, no additional policy is applied for greylist user. Greylist user will get access permissions same as that of actual user logged in.

When user deletes an entry from greylist, identity manager will:

1. Delete the entry and updates the list.
2. User identity is constructed based on *NetLogin* details, if deleted username is found in NetLogin authenticated user database.

Example

The following command adds an username to the greylist:

```
configure identity-management greylist add user Richard@corp
```

The following command deletes an username from the greylist:

```
configure identity-management greylist del user Richard@corp
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure identity-management kerberos snooping aging time

```
configure identity-management kerberos snooping aging time minutes
```

Description

Specifies the aging time for Kerberos snooping entries.

Syntax Description

<i>minutes</i>	Specifies the aging time in minutes. The range is 1 to 65535 minutes.
----------------	---

Default

N/A.

Usage Guidelines

Kerberos does not provide any service for un-authentication or logout. Kerberos does provide a ticket lifetime, but that value is encrypted and cannot be detected during snooping.

To enable the aging and removal of snooped Kerberos entries, this timer defines a maximum age for the snooped entry. When a MAC address with a corresponding Kerberos entry in Identity Manager is aged out, the Kerberos snooping timer starts. If the MAC address becomes active before the Kerberos snooping timer expires, the timer is reset and the Kerberos entry remains active. If the MAC address is inactive when the Kerberos snooping timer expires, the Kerberos entry is removed.

Example

The following command configures the aging time for 600 minutes:

```
* Switch.4 # configure identity-management kerberos snooping aging time 600
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure identity-management kerberos snooping force-aging time

```
configure identity-management kerberos snooping force-aging time [none |  
minutes]
```

Description

Configures the switch to remove all Kerberos snooping entries after the specified time expires.

Syntax Description

<i>minutes</i>	Specifies the aging time in minutes. The range is 1 to 65535 minutes.
none	Disables the Kerberos force-aging feature.

Default

N/A.

Usage Guidelines

If Kerberos force aging is enabled, we recommend that the Kerberos snooping force aging time be set to the same value as the Kerberos ticket lifetime.

Example

The following command removes all Kerberos snooping entries after 600 minutes:

```
* Switch.4 # configure identity-management kerberos snooping force-aging time 600
```

History

This command was first available in ExtremeXOS 12.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure identity-management kerberos snooping forwarding

```
configure identity-management kerberos snooping forwarding [fast-path | slow-path]
```

Description

When identity management is enabled on a port, Kerberos packets are software-forwarded. With this command, you can report if shared folder access via identity management-enabled ports is slow if there exists other CPU-bound traffic.

Syntax Description

forwarding	Configure how customer Kerberos authentication packets are forwarded by this system.
fast-path	Forward customer snooped Kerberos packets in hardware (default).
slow-path	Forward customer snooped Kerberos packets in software. This option is recommended only for systems with low CPU-bound traffic.

Default

Fast-path.

Usage Guidelines

Use this command to report if shared folder access via identity management-enabled ports is slow if there exists other CPU-bound traffic.

Example

The following show command displays the modified Kerberos information:

```
# sh identity-management
Identity Management           : Enabled
Stale entry age out (effective) : 180 Seconds (180 Seconds)
Max memory size              : 512 Kbytes
Enabled ports                 : 1
SNMP trap notification       : Enabled
Access list source address type : MAC
Kerberos aging time (DD:HH:MM) : None
Kerberos force aging time (DD:HH:MM) : None
Kerberos snooping forwarding  : Fast path
Kerberos snooping forwarding  : Slow path
Valid Kerberos servers        : none configured(all valid)
LDAP Configuration:
-----
LDAP Server      : No LDAP Servers configured
Base-DN          : None
Bind credential  : anonymous

LDAP Configuration for Netlogin:
dot1x            : Enabled
mac              : Enabled
web-based        : Enabled
```

History

This command was first available in ExtremeXOS 15.1.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure identity-management kerberos snooping server

```
configure identity-management kerberos snooping add server ip_address
configure identity-management kerberos snooping delete server
                               [ip_address | all]
```

Description

Adds or deletes a Kerberos server to the Kerberos server list.

Syntax Description

<i>ip_address</i>	Specifies a Kerberos server IP address to add or delete.
all	Specifies that all Kerberos server list entries are to be deleted.

Default

No servers are in the Kerberos server list.

Usage Guidelines

When no servers are configured in the Kerberos server list, the Kerberos snooping feature processes responses from all Kerberos servers, which can expose the system to simulated logins. To avoid this exposure, you can configure a list of up to 20 valid Kerberos servers. When the Kerberos server list contains one or more entries, the switch only processes responses from the Kerberos servers in the list.

Example

The following command adds the Kerberos server at IP address 10.10.10.1 to the Kerberos server list:

```
* Switch.4 # configure identity-management kerberos snooping add server 10.10.10.1
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure identity-management list-precedence

```
configure identity-management list-precedence listname1 listname2  
listname3
```

Description

This command allows you to configure the precedence of list types. You must specify the list-names in the desired order of precedence. Listname1 will take precedence of all lists (i.e., highest precedence). Listname2 will take precedence over Listname3. When the user/device logs in, entries present in Listname1 will be searched at first to find matching role. Entries present in Listname2 will be searched after Listname1 and entries in Listname3 will be searched at last.

Syntax Description

<i>listname1</i>	Specifies the list type which has precedence over all list types.
<i>listname2</i>	Specifies the list type which has next precedence, after listname1.
<i>listname3</i>	Specifies the list type which has least precedence of all.

Default

greylist, blacklist, whitelist

Usage Guidelines

By default, greylist entries have higher precedence over blacklist and whitelist entries.

This means that IDM consults with greylist first upon detection of user, and then decides if identity needs to be created. If there is a greylist entry matching the incoming username, user identity is not created. If there is no matching greylist entry, IDM proceeds with role identification for the user. However, greylist precedence is configurable. Following are three possibilities for greylist precedence configuration.

1. greylist, blacklist, whitelist
2. blacklist, greylist, whitelist
3. blacklist, whitelist, greylist

It is important to notice that blacklist always has higher precedence over whitelist for ExtremeXOS 15.1.2. In order to change the list precedence, Identity Management should be disabled first. Disabling IDM is required since there may be many users/devices already mapped to some roles and policies/ACLs applied. Considering the processing load of unmapping the roles and removing policies, changing precedence isn't allowed when IDM is enabled. When precedence configuration is changed, each entry present in the list with lower precedence (new precedence) is checked with each entry present in all the lists with higher precedence.

Example

The following example instructs that blacklist has precedence over all lists. Greylist has precedence over whitelist. Whitelist has least precedence.

```
configure identity-management list-precedence blacklist greylist whitelist
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure identity-management ports

```
configure identity-management {add | delete} ports [port_list | all]
```

Description

Adds or deletes identity management for the specified ports.

Syntax Description

add	Enables identity management on the specified port list.
delete	Disables identity management on the specified port list.
<i>port_list</i>	Specifies the ports to which this command applies.
all	Specifies that this command applies to all ports.

Default

No ports are in the identity management enabled port list.

Usage Guidelines

If neither the add nor the delete keyword is entered, identity management is enabled on the specified port list, and the new port list overrides any previous port list.

If identity management is enabled on a port and a user or device is connected to it, information about the user or device is present in the identity management database. If this port is removed from the identity-management enabled port list, the user or device information remains in the data base until the user logs out or the device disconnects. However, once a port is deleted from enabled port list, no new information is added to the identity management database for that port.



Note

Kerberos identities are not detected when both server and client ports are added to identity management.

Example

The following command enables identity management on ports 2:3 and 2:5:

```
configure identity-management add ports 2:3,2:5
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure identity-management role add child-role

```
configure identity-management role role_name add child-role child_role
```

Description

Adds a child role to the specified role.

Syntax Description

<i>role_name</i>	Specifies the name of an existing role.
child-role	Specifies a name for the new child role (up to 32 characters).

Default

N/A.

Usage Guidelines

The child role name can include up to 32 characters. Role names must begin with an alphabetical letter, and only alphanumeric, underscore (_), and hyphen (-) characters are allowed in the remainder of the name. Role names cannot match reserved keywords. For more information on role name requirements and a list of reserved keywords, see [Object Names](#) on page 11.

The following guidelines apply to child roles:

- A child role inherits all the policies applied to its parent and any higher levels above the parent.
- The software supports 5 levels of hierarchy.
- Each role can have a maximum of 8 child roles.
- Each child role can have only 1 parent role.

Example

The following example configures a child role named East for the existing role named India-Engr:

```
* Switch.66 # configure identity-management role "India-Engr" add child-role East
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure identity-management role add dynamic-rule

```
configure identity-management role role_name [add dynamic-rule rule_name
{ first | last | { [before | after] ref_rule_name } }]
```

Description

Adds a dynamic ACL rule for the specified role and specifies the order.

Syntax Description

<i>role_name</i>	Specifies the name of an existing role.
<i>rule_name</i>	Specifies the name of a dynamic ACL rule to add to the specified role.

Default

The order of the dynamic rule is last if the order is not explicitly specified.

Usage Guidelines

The maximum number of policies or ACL rules that can be applied to a particular role is restricted to 8. This count does not include the policies and rules inherited from a parent role. Since the maximum hierarchy depth is 5, the maximum number of policies and rules supported for a role at the maximum hierarchy depth is 40 (8 x 5).

When a dynamic ACL rule is added to a role, it is immediately installed for all identities mapped to that role and roles below it in the role hierarchy.

Example

The following example configures the role named India-Engr to use the ACL rule named india-Engr-rule:

```
* Switch.55 # configure identity-management role "India-Engr" add dynamic-rule india-Engr-rule
```

History

This command was first available in ExtremeXOS 12.5.

This command was modified in ExtremeXOS 15.2.1 to specify order.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure identity-management role add policy

```
configure identity-management role role_name add policy policy-name
    {first | last {[before | after] ref_policy_name}}
```

Description

Adds a policy for the specified role and specifies the order.

Syntax Description

<i>role_name</i>	Specifies the name of an existing role.
<i>policy-name</i>	Specifies the name of a policy to add to the specified role.

Default

The order of the policy is last if the order is not explicitly specified.

Usage Guidelines

The maximum number of policies or [ACL](#) rules that can be applied to a particular role is restricted to 8. This count does not include the policies and rules inherited from a parent role. Since the maximum hierarchy depth is 5, the maximum number of policies and rules supported for a role at the maximum hierarchy depth is 40 (8 x 5).

When a policy is added to a role, it is immediately installed for all identities mapped to that role and all roles below it in the role hierarchy.

Example

The following example configures the role named India-Engr to use the policy named india-Engr-policy:

```
* Switch.44 # configure identity-management role "India-Engr" add policy india-Engr-policy
```

History

This command was first available in ExtremeXOS 12.5.

This command was modified in ExtremeXOS 15.2.1 to specify order.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure identity-management role delete child-role

```
configure identity-management role role_name delete child-role
[child_role | all]
```

Description

Deletes one or all child roles from the specified role.

Syntax Description

<i>role_name</i>	Specifies the name of an existing role.
child-role	Specifies a name for a child role to delete.
all	Specifies that all child roles are to be deleted.

Default

N/A.

Usage Guidelines

None.

Example

The following example deletes the child role named East from the existing role named India-Engr:

```
* Switch.66 # configure identity-management role "India-Engr" delete child-role East
```

The following command deletes all child roles from the existing role named India-Engr:

```
* Switch.66 # configure identity-management role "India-Engr" delete child-role all
```

History

This command was first available in ExtremeXOS 12.5.

The all option was added in ExtremeXOS 12.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure identity-management role delete dynamic-rule

```
configure identity-management role role_name delete dynamic-rule
[rule_name | all]
```

Description

Deletes one or all dynamic *ACL* rules for the specified role.

Syntax Description

<i>role_name</i>	Specifies the name of an existing role.
<i>rule_name</i>	Specifies the name of a dynamic ACL rule to delete from the specified role.
all	Specifies that all dynamic ACL rules are to be deleted.

Default

N/A.

Usage Guidelines

None.

Example

The following example deletes all dynamic rules from the role named India-Engr:

```
* Switch.55 # configure identity-management role "India-Engr" delete dynamic-rule all
```

History

This command was first available in ExtremeXOS 12.5.

The all option was added in ExtremeXOS 12.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure identity-management role delete policy

```
configure identity-management role role_name delete policy [policy-name  
| all]
```

Description

Deletes one or all policies for the specified role.

Syntax Description

<i>role_name</i>	Specifies the name of an existing role.
<i>policy-name</i>	Specifies the name of a policy to delete from the specified role.
all	Specifies that all policies are to be deleted from the specified role.

Default

N/A.

Usage Guidelines

None.

Example

The following example deletes the policy named india-Engr-policy from the role named India-Engr:

```
* Switch.44 # configure identity-management role "India-Engr" delete policy india-Engr-policy
```

History

This command was first available in ExtremeXOS 12.5.

The all option was added in ExtremeXOS 12.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure identity-management role match-criteria inheritance

```
configure identity-management role match-criteria inheritance [on | off]
```

Description

This command enables or disables the match-criteria inheritance support. Check the current status by issuing the `show identity-management` command.

Syntax Description

<i>role</i>	User role.
<i>match-criteria</i>	Match criteria for the role.
<i>inheritance</i>	Inheriting match criteria from parent role to child role.
on off	Specifies whether match criteria inheritance is on or off.

Default

Off.

Usage Guidelines

From ExtremeXOS Release 15.2, child roles can inherit the match criteria of the parent role. This helps the user since the match criteria need not be duplicated in all levels of hierarchy.

When match-criteria inheritance is on, for a user to be classified under a child role, he has to satisfy the match criteria of the child role, and also all parent roles in the hierarchy.

Match criteria inheritance helps users in avoiding the need to duplicate match-criteria entries in the hierarchy.

Example

For example, there are roles called Employee, USEmployee and USSales in an organization hierarchy of a company XYZCorp.com. Till ExtremeXOS 15.1 (or with match-criteria inheritance off), the user has to create three roles like this:

```
* Switch.1 # create identity-management role Employee match-criteria "company ==
XYZCorp.com;"
* Switch.2 # create identity-management role USEmployee match-criteria "company ==
XYZCorp.com; AND country == USA;"
* Switch.3 # create identity-management role USSales match-criteria "company ==
XYZCorp.com; AND country == USA; AND department = Sales"
* Switch.4 # configure identity-management role "Employee" add child-role "USEmployee"
* Switch.5 # configure identity-management role "USEmployee" add child-role "USSales"
```

Now this can be simplified into the following since child role inherits parent role's match criteria:

```
* Switch.1 # configure identity-management role match-criteria inheritance on
* Switch.2 # create identity-management role Employee match-criteria "company ==
XYZCorp.com;"
* Switch.3 # create identity-management role USEmployee match-criteria "country == USA;"
* Switch.4 # create identity-management role USSales match-criteria "department = Sales"
* Switch.5 # configure identity-management role "Employee" add child-role "USEmployee"
* Switch.6 # configure identity-management role "USEmployee" add child-role "USSales"
```

History

This command was first available in ExtremeXOS 15.2

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure identity-management role priority

```
configure identity-management role role_name priority pri_value
```

Description

Configures a priority value for the specified role.

Syntax Description

<i>role_name</i>	Specifies the name of an existing role that you want to configure.
<i>pri_value</i>	Specifies the role priority; the lower the priority number, the higher the priority. The range of values is 1 to 255. Value 1 represents the highest priority, and value 255 represents the lowest priority.

Default

Priority=255.

Usage Guidelines

The role priority determines which role a user is mapped to when the user's attributes match the match-criteria of more than 1 role. If the user's attributes match multiple roles, the highest priority (lowest priority value) role applies. If the priority is the same for all matching roles, the role for which the priority was most recently set or modified is used.

Example

The following example configures the role named India-Engr to use the highest priority:

```
* Switch.33 # configure identity-management role "India-Engr" priority 1
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure identity-management stale-entry aging-time

```
configure identity-management stale-entry aging-time seconds
```

Description

Configures the stale-entry aging time for event entries in the identity management database.

Syntax Description

<i>seconds</i>	Specifies the period (in seconds) at which event entries are deleted. The range is 60 to 1800 seconds.
----------------	--

Default

180 seconds.

Usage Guidelines

The identity management database contains active entries, which correspond to active users and devices, and event entries, which record identity management events such as user logout or device disconnect. The active entries are automatically removed when a user logs out or a device disconnects. The event entries are automatically removed after a period defined by the stale-entry aging time.



Note

To capture active and event entries before they are deleted, you can use external management software such as Ridgeline™, which can access the switch using XML APIs. We recommend that the external client(s) that poll the identity management database be configured for polling cycles that are between one-third and two-thirds of the stale-aging time. This ensures that a new database entry or event does not age out before the next polling cycle.

The stale-entry aging time defines when event entries become stale. To preserve memory, the software periodically uses a cleanup process to remove the stale entries. You can configure the stale-entry aging time. The cleanup interval is defined by the software.

When memory usage is high, the software reduces both the stale-entry aging time and the cleanup interval to keep memory available for new entries. The following table shows how the database is managed as memory usage increases.

Table 8: Identity Management Database Usage Levels

Database Memory Usage Level	Database Memory Usage Level (Percent)	Effective Stale-Entry Aging Time	Description
Normal	Up to 80%	Configured stale-entry aging time	New identities and associated information (VLAN and IP addresses) are added to or updated in the database. Events are also added to the database. Events are deleted from the database after the configured stale-entry aging time.
High	Above 80% to 90%	The lower value of the following: 90 seconds or 50% of the configured stale-entry aging time	Identities and events are added to the database as for the normal usage level, but the effective stale-entry aging time is reduced to delete events sooner and free memory.

Table 8: Identity Management Database Usage Levels (continued)

Database Memory Usage Level	Database Memory Usage Level (Percent)	Effective Stale-Entry Aging Time	Description
Critical	Above 90%	15 seconds	The effective stale-entry aging time is further reduced to delete events sooner and free memory. No new identities are added to the database at this usage level, but updates (such as the addition or deletion of a VLAN or IP address) continue. At this level, the database might be missing active entries.
Maximum	Above 98%	15 seconds	At this level, the software does not process additions or updates to the database. The software only processes deletions. At this level, the database might be missing active entries.

Whenever the database usage level changes, an EMS message is logged, and if enabled, an *SNMP* trap is sent. If the switch changes the stale-entry aging time, the SNMP trap contains the new stale-entry aging time.

**Note**

If the database level regularly reaches the high usage level, or if it reaches the critical or maximum levels, it is time to investigate the cause of the issue. The solution might be to increase the database memory size.

External clients should be capable of adjusting the polling cycles. Because the aging cycle is shorter when memory is low, it is best if external clients can adjust their polling cycles in response to SNMP traps that announce a change in the stale-entry aging time.

Example

The following command configures the stale-entry aging time for 90 seconds:

```
* Switch.4 # configure identity-management stale-entry aging-time 90
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure identity-management whitelist

```
configure identity-management whitelist add [mac mac_address {macmask}
| ip ip_address {netmask} | ipNetmask] | user user_name] configure
identity-management whitelist delete [all | mac mac_address {macmask}
| ip ip_address {netmask} | ipNetmask] | user user_name]
```

Description

Adds or deletes an identity in the identity manager whitelist.

Syntax Description

add	Adds the specified identity to the whitelist.
delete	Deletes the specified identity from the whitelist.
all	Specifies that all identities are to be deleted from the whitelist. This option is available only when the delete attribute is specified.
<i>mac_address</i>	Specifies an identity by MAC address.
<i>macmask</i>	Specifies a MAC address mask. For example: FF:FF:FF:00:00:00.
<i>ip_address</i>	Specifies an identity by IP address.
<i>netmask</i>	Specifies a mask for the specified IP address.
<i>ipNetmask</i>	Specifies an IP network mask.
<i>user_name</i>	Specifies an identity by user name.

Default

N/A.

Usage Guidelines

The software supports up to 512 entries in the whitelist. When you add an identity to the whitelist, the switch searches the blacklist for the same identity. If the identity is already in the blacklist, the switch displays an error.

It is possible to configure an identity in both lists by specifying different attributes in each list. For example, you can add an identity username to the whitelist and add the MAC address for that user's laptop in the blacklist. Because the blacklist has priority over the whitelist, identity access is denied from the user's laptop, but the user can access the switch from other locations.

If you add a new whitelist entry that is qualified by a MAC or IP address, the identity manager does the following:

- Reviews the identities already known to the switch. If the new whitelist entry is blacklisted (by specifying a different identity attribute), no action is taken.
- If the identity is not blacklisted and is known on the switch, all existing [ACLs](#) for the identity are removed.

- When a whitelisted MAC-based identity is detected or already known, an Allow All ACL is programmed for the identity MAC address for the port on which the identity is detected.
- When a whitelisted IP-based identity is detected or already known, an Allow All ACL is programmed for the identity IP address for the port on which the identity is detected.

If you add a new whitelist entry that is qualified by a username (with or without a domain name), the identity manager does the following:

- Reviews the identities already known to the switch. If the new whitelist entry is an identity known on the switch, an Allow All ACL is programmed for the identity MAC address on all ports to which the identity is connected.
- When a new whitelisted username-based identity accesses the switch, an Allow All ACL is programmed for the identity MAC address on the port on which the identity is detected.
- The ACL for a whitelisted username follows any ACLs based on Kerberos snooping.

Allow All ACLs for whitelisted entries exist as long as the identity remains in the identity manager database.

If you delete an identity from the whitelist, identity manager checks to see if the identity is in the local database. If the identity is known to the switch, the switch does the following:

- Removes the Allow All ACL from the port to which the identity connected.
- Initiates the role determination procedure for the switch port to which the known identity connected. This ensures that the appropriate role is applied to the identity that is no longer whitelisted.



Note

The role determination process can trigger an LDAP refresh to collect identity attributes for role determination.

Example

The following command adds an IP address to the whitelist:

```
* Switch.4 # configure identity-management whitelist add ip 10.0.0.1
```

The following command deletes a user name from the whitelist:

```
* Switch.5 # configure identity-management whitelist delete user john
```

History

This command was first available in ExtremeXOS 12.7.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure cli idle-timeout

```
configure cli idle-timeout minutes
```

Description

Configures the time-out for idle console, SSH2, and Telnet sessions.

Syntax Description

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

Default

The default time-out is 20 minutes.

Usage Guidelines

This command configures the length of time the switch will wait before disconnecting idle console, SSH2, or Telnet sessions.

The `idletimeout` feature must be enabled for this command to have an effect (the `idletimeout` feature is enabled by default).

Example

The following command sets the time-out for idle login and console sessions to 10 minutes:

```
configure cli idle-timeout 10
```

History

This command was first available in ExtremeXOS 10.1.

The **cli** keyword was added and the **idletimeout** keyword was changed to **idle-timeout** in ExtremeXOS 30.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure igmp

```
configure igmp query_interval query_response_interval  
               last_member_query_interval {{vlan} vlan_name} {{vr} vr_name}  
               {{robustness}}
```


Description

Configures the *IGMP* timers.

Syntax Description

<i>query_interval</i>	Specifies the interval (in seconds) between general queries.
<i>query_response_interval</i>	Specifies the maximum query response time (in seconds).
<i>last_member_query_interval</i>	Specifies the maximum group-specific query response time (in seconds).
<i>vlan_name</i>	Applies the configuration only to the specified <i>VLAN</i> . If no VLAN is specified, the configuration applies to all VLANs.
<i>vr_name</i>	Specifies the VR to which the configuration should be applied. If no parameter is specified, the configuration is applied to the current VR context.
<i>robustness</i>	Specifies the degree of robustness for the network.

Default

- query interval—125 seconds
- query response interval—10 seconds
- last member query interval—1 second
- robustness—2

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.
- robustness—The degree of robustness of the network. The range is 2 to 7. This parameter allows tuning for the expected packet loss on a link. If a link is expected to have packet loss, this parameter can be increased.
- The group timeout is defined by the formula: $\text{group_timeout} = (\text{query_interval} \times \text{robustness}) + \text{query_response_interval}$, according to RFC 2236. You can explicitly define the host timeout using the `configure igmp snooping timer router_timeout host_timeout {vr vrname}` command. The effective *host_timeout* is the lesser value of the *group_timeout* and the configured *host_timeout*.

Example

The following command configures the IGMP timers:

```
configure igmp 100 5 1 3
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IPv4 multicast feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure igmp router-alert receive-required

```
configure igmp router-alert receive-required [on | off] {{vlan}
vlan_name}
```

Description

Controls when the router-alert option is required for IGMPv2 and IGMPv3 packet reception and processing.

Syntax Description

vlan	Applies the configuration only to the specified <u>VLAN</u> . If no VLAN is specified, the configuration applies to all VLANs.
-------------	--

Default

Off—All IGMP packets are received and processed.

Usage Guidelines

By default, the ExtremeXOS software receives and processes all IGMP packets, regardless of the setting of the router-alert option within a packet. The default configuration works with all switches that support the ExtremeXOS software.

IETF standards require that a router accept and process IGMPv2 and IGMPv3 packets only when the router-alert option is set. The on setting for this command sets the ExtremeXOS software to comply with the IETF standards and should be used when the switch will be used with third-party switches that expect IETF compliant behavior.

Example

The following command configures the switch for IETF compliant IGMP packet processing:

```
configure igmp router-alert receive-required on
```

History

This command was first available in ExtremeXOS 12.5.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IPv4 multicast feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure igmp router-alert transmit

```
configure igmp router-alert transmit [on | off] [{vlan} vlan_name]
```

Description

Controls whether the **router-alert** option is set when forwarding IGMPv2 and IGMPv3 packets.

Syntax Description

vlan	Applies the configuration only to the specified <i>VLAN</i> . If no VLAN is specified, the configuration applies to all VLANs.
-------------	--

Default

On—The **router-alert** option is set when forwarding IGMPv2 and IGMPv3 packets.

Usage Guidelines

IETF standards require that a router set the **router-alert** option in forwarded IGMPv2 and IGMPv3 packets. The ExtremeXOS software has been updated to comply with this requirement using the default settings.

Earlier versions of the ExtremeXOS software forwarded all *IGMP* packets without setting the **router-alert** option. If compatibility issues arise, you can configure the software to use the legacy behavior by using this command with the **off** option.

Example

The following command configures the switch for IETF compliant IGMP packet processing:

```
configure igmp router-alert transmit on
```

History

This command was first available in ExtremeXOS 12.5.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IPv4 multicast feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure igmp snooping filters

```
configure igmp snooping filters [per-port | per-vlan]
```

Description

Selects the type of *IGMP* snooping filters that are installed.

Syntax Description

per-port	Installs the per-port IGMP snooping filters.
per-vlan	Installs the per- <i>VLAN</i> IGMP snooping filters.

Default

per-port.

Usage Guidelines

Use the per-vlan option when the number of VLANs configured on the switch is lower than the maximum numbers listed in the following table. This option conserves usage of the hardware Layer 3 multicast forwarding table.

When the number of configured VLANs is larger than the maximum values listed here, select the per-port option. Each VLAN requires additional interface hardware *ACL* resources. The per-port option conserves usage of the interface hardware ACL resources.

Table 9: Maximum Number of VLANs Supported by per-VLAN IGMP Snooping Filters

ExtremeSwitching Switch Series Module Type	Maximum Number of VLANs When per-VLAN Snooping Filters are Installed
a Series	1000
c Series	2000.
e Series	448.
xl Series	2000.

The actual maximum value is smaller if other processes require entries in the interface ACL table. To display the IGMP snooping filters configuration, use the `show igmp snooping` command.



Note

For MLD Snooping, the maximum number of VLANs is half of the numbers provided in this table. The maximum number specified here is individual limit for IGMP snooping filters. If both IGMP and MLD snooping filters are used, the maximum numbers are lower than the ones specified.

Example

The following command configures the switch to install the per-VLAN IGMP snooping filters:

```
# configure igmp snooping filters per-vlan
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IGMP snooping feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure igmp snooping flood-list

```
configure igmp snooping flood-list [policy | none] {vr vrname}
```

Description

Configures certain multicast addresses to be slow path flooded within the VLAN.

Syntax Description

<i>policy</i>	Specifies a policy file with a list of multicast addresses to be handled.
none	Specifies no policy file is to be used.
<i>vrname</i>	Specifies a virtual router.

Default

None.

Usage Guidelines

With this command, a user can configure certain multicast addresses to be slow path flooded within the VLAN, which otherwise are fast path forwarded according to IGMP and/or Layer 3 multicast protocol.

A policy file is a text file with the extension, .pol. It can be created or edited with any text editor. The specified policy file *policy file* should contain a list of addresses which determine if certain multicast streams are to be treated specially. Typically, if the switch receives a stream with a destination address which is in the *policy file* in 'permit' mode, that stream is software flooded and no hardware entry is installed.

When adding an IP address into the policy file, 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 streams as control packets.

To create a policy file for the snooping flood-list, use the following template:

```
# This is a template for IGMP Snooping Flood-list Policy File
# Add your group addresses between "Start" and "End"
# Do not touch the rest of the file!!!
entry igmpFlood {
    if match any {
#----- Start of group addresses -----
        nlri 234.1.1.1/32;
        nlri 239.1.1.1/32;
#----- end of group addresses -----
    }
    then {
        permit;
    }
}
entry catch_all {
    if {
    }
    then {
        deny;
    }
}
```



Note

The switch does not validate any IP address in the policy file 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 is done within the L2 VLAN only.

Use the **none** option to effectively disable slow path flooding.

You can use the `show igmp` command to see the configuration of slow path flooding.

Example

The following example configures the multicast data stream specified in access1 for slow path flooding:

```
configure igmp snooping flood-list access1
```

The following command specifies that no policy file is to be used, this effectively disabling slow path flooding:

```
configure igmp snooping flood-list none
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IGMP snooping feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure igmp snooping leave-timeout

```
configure igmp snooping leave-timeout leave_timeout_ms [{vlan}  
  vlan_name] [{vr} vr_name]
```

Description

Configures the *IGMP* snooping leave timeout.

Syntax Description

<i>leave_timeout_ms</i>	Specifies an IGMP leave timeout value in milliseconds.
<i>vlan_name</i>	Applies the configuration only to the specified <i>VLAN</i> . If no VLAN is specified, the configuration applies to all VLANs.

Default

1000 ms.

Usage Guidelines

The leave-timeout is the IGMP leave override interval. If no other hosts override the IGMP leave by the end of this interval, the receiver port is removed.

The range is 0 - 175000 ms (175 seconds). For timeout values of one second or less, you must set the leave-timeout to a multiple of 100 ms. For values of more than one second, you must set the leave-timeout to a multiple of 1000 ms (one second).

Example

The following example configures the IGMP snooping leave timeout to one second:

```
configure igmp snooping leave-timeout 1000
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IGMP snooping feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure igmp snooping timer

```
configure igmp snooping timer router_timeout host_timeout {vr vrname}
                        {vlan vlan_name}
```

Description

Configures the [IGMP](#) snooping timers.

Syntax Description

<i>router_timeout</i>	Specifies the time in seconds before removing a router snooping entry.
<i>host_timeout</i>	Specifies the time in seconds before removing a host's group snooping entry.
<i>vrname</i>	Specifies a virtual router.
<i>vlan_name</i>	Specifies the VLAN name. If no VLAN is specified, the setting is applied to all existing VLANs.

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 maximum time, in seconds, that a router snooping entry can remain in the IGMP snooping table without receiving a router report. If a report is not received, the entry is deleted. The range is 10 to 214,748,364 seconds (6.8 years). The default setting is 260 seconds.
- *host_timeout*—The maximum time, in seconds, that a group snooping entry can remain in the IGMP snooping table without receiving a group report. If a report is not received, the entry is deleted. The range is 10 to 214,748,364 seconds. The default setting is 260 seconds.



Note

The *host_timeout* value should be less than or equal to the query timeout value, which is defined by the following: (query_interval x robustness) + query_response_interval.

IGMP snooping expects at least one device on every VLAN to periodically generate IGMP query messages. Without an IGMP querier, the switch eventually stops forwarding IP multicast packets to any port, because the IGMP snooping entries time out, based on the value specified in *host_timeout* or *router_timeout*.

Example

The following example configures the IGMP snooping timers:

```
configure igmp snooping timer 600 600
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IGMP snooping feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure igmp snooping vlan ports add dynamic group

```
configure igmp snooping {vlan} vlan_name {ports portlist} add dynamic
group [ grpipaddress ]
```

Description

Configures an *IGMP* dynamic group.

Syntax Description

<i>vlan_name</i>	Specifies a vlan name.
<i>portlist</i>	Specifies a port list.
<i>grpipaddress</i>	Specifies the multicast group IP address.

Default

N/A.

Usage Guidelines

This command adds IGMP groups to specific VLANs or to ports belonging to specific VLANs. After the groups are added, the expiration timer is started. This causes the groups to expire. The configuration is not saved in the configuration file. The following message is displayed on execution of this command:

```
INFO: This command is not saved in the configuration.
```

Example

The following example adds a dynamic group to a switch port:

```
switch.111 # configure igmp snooping vlan "ixia113" ports 47 add dynamic group 225.1.1.1
INFO: This command is not saved in the configuration.
```

The following command displays the group:

```
switch.112 # show igmp group
Group Address      Ver Vlan      Port      Age
225.1.1.1          2   ixia113    47         3

Total: 1
switch.113 #
```

The following example adds a dynamic group to a vlan (loopback port):

```
switch sw5.113 # configure igmp snooping vlan "ixia113" add dynamic group 225.1.1.1
INFO: This command is not saved in the configuration.
```

The following command displays the group:

```
switch.114 # show igmp group
Group Address      Ver Vlan      Port      Age
225.1.1.1          2   ixia113    Lpbk       37

Total: 1
switch.115 #
```

History

This command was first available in ExtremeXOS 15.3.2.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IPv4 multicast feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure igmp snooping vlan ports add static group

```
configure igmp snooping {vlan} vlanname {ports portlist }add static
group grpipaddress
```

Description

Configures VLAN ports to receive the traffic from a multicast group, even if no IGMP joins have been received on the port.

Syntax Description

<i>vlanname</i>	Specifies a VLAN name.
<i>portlist</i>	Specifies one or more ports or slots and ports.
<i>grpipaddress</i>	Specifies the multicast group IP address.

Default

N/A.

Usage Guidelines

Use this command to forward a particular multicast group to VLAN ports. In effect, this command emulates a host on the port that has joined the multicast group. As long as the port is configured with the static entry, multicast traffic for that multicast group is forwarded to that port.

This command is for IGMPv2 only.

The switch sends proxy IGMP messages in place of those generated by a real host. The proxy messages use the VLAN IP address for source address of the messages. If the VLAN has no IP address assigned, the proxy IGMP message uses 0.0.0.0 as the source IP address.

The multicast group should be in the class-D multicast address space, but should not be in the multicast control subnet range (224.0.0.x/24).

If the ports also have an IGMP filter configured, the filter entries take precedence. IGMP filters are configured using the command:

```
configure igmp snooping vlan vlanname ports portlist filterpolicy file
```

Example

The following example configures a static IGMP entry so that multicast group 225.1.1.1 is forwarded to VLAN "marketing" on port 47:

```
switch.30 # configure igmp snooping marketing ports 47 add static group 225.1.1.1
```

The following command displays the group:

```
* (pacman debug) sw4.31 # show igmp group
Group Address      Ver Vlan      Port    Age
225.1.1.1(s)      2   marketing  47      0

Total: 1
switch.32 #
```

The following example adds a static group to a vlan (loopback port):

```
switch.32 # configure igmp snooping marketing add static group 225.1.1.1
```

The following command displays the group:

```
switch.33 # show igmp group
Group Address      Ver Vlan      Port    Age
225.1.1.1(s)      2   marketing  Lpbk    0
```

```
Total: 1
switch.34 #
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IGMP snooping feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure igmp snooping vlan ports add static router

```
configure igmp snooping {vlan} vlanname ports portlist add static router
```

Description

Configures VLAN ports to forward the traffic from all multicast groups, even if no IGMP joins have been received on the port.

Syntax Description

<i>vlanname</i>	Specifies a VLAN name.
<i>portlist</i>	Specifies one or more ports or slots and ports.

Default

N/A.

Usage Guidelines

Use this command to forward all multicast groups to the specified VLAN ports. In effect, this command emulates a multicast router attached to those ports. As long as the ports are configured with the static entry, all available multicast traffic is forwarded to those ports.

Example

The following example configures a static IGMP entry so all multicast groups are forwarded to VLAN marketing on ports 2:1-2:4:

```
configure igmp snooping marketing ports 2:1-2:4 add static router
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IGMP snooping feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure igmp snooping vlan ports delete static group

```
configure igmp snooping {vlan} vlan_name {ports port_list} delete static
group [ip_address | all]
```

Description

Removes the port configuration that causes multicast group traffic to be forwarded, even if no *IGMP* leaves have been received on the port.

Syntax Description

<i>vlan_name</i>	Specifies a <i>VLAN</i> name.
<i>port_list</i>	Specifies one or more ports or slots and ports.
<i>ip_address</i>	Specifies the multicast group IP address.
all	Delete all the static groups.

Default

N/A.

Usage Guidelines

This command is used to remove a static IGMP group entry created on a VLAN or on a port. Use this command to remove a static group entry created by the following command:

```
configure igmp snooping vlan vlanname ports portlist add static group
ipaddress
```

Example

The following example removes a static IGMP entry that forwards the multicast group 224.34.15.37 to the VLAN marketing on ports 2:1-2:4:

```
configure igmp snooping marketing ports 2:1-2:4 delete static group 224.34.15.37
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IGMP snooping feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure igmp snooping vlan ports delete static router

```
configure igmp snooping vlan vlanname ports portlist delete static  
router
```

Description

Removes the configuration that causes VLAN ports to forward the traffic from all multicast groups, even if no IGMP joins have been received on the port.

Syntax Description

<i>vlanname</i>	Specifies a VLAN name.
<i>portlist</i>	Specifies one or more ports or slots and ports. On a SummitStack, it can be a list of slots and ports. On a standalone 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 is used to remove a static router port entry created on a VLAN. Use this command to remove an entry created by the following command:

```
configure igmp snooping vlan vlanname ports portlist add static router
```

Example

The following example removes the static IGMP entry that caused all multicast groups to be forwarded to VLAN marketing on ports 2:1-2:4:

```
configure igmp snooping marketing ports 2:1-2:4 delete static router
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IGMP snooping feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure igmp snooping vlan ports filter

```
configure igmp snooping vlan vlanname ports portlist filter [policy |  
none]
```

Description

Configures an IGMP snooping policy file filter on VLAN ports.

Syntax Description

<i>vlanname</i>	Specifies a VLAN name.
<i>portlist</i>	Specifies one or more ports or slots and ports. On a SummitStack, it can be a list of slots and ports. On a stand-alone switch, it can be one or more port numbers. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.
<i>policy</i>	Specifies the policy file for the filter.

Default

None.

Usage Guidelines

Use this command to filter multicast groups to the specified VLAN ports.

The policy file used by this command is a text file that contains the class-D addresses of the multicast groups that you wish to block.

To remove IGMP snooping filtering from a port, use the **none** keyword version of the command.

Use the following template to create a snooping filter policy file:

```
# # Add your group addresses between "Start" and "end" # Do not touch the rest of the  
file!!!!  
entry igmpFilter  
{ if match any  
{  
#----- Start of group addresses -----  
nlri 239.11.0.0/16; nlri 239.10.10.4/32;  
#----- end of group addresses -----  
} then { deny;  
}  
}  
entry catch_all  
{ if  
{  
} then
```

```
{ permit;  
}  
}
```

Example

The following example configures the policy file `ap_multicast` to filter multicast packets forwarded to VLAN marketing on ports 2:1-2:4:

```
configure igmp snooping marketing ports 2:1-2:4 filter ap_multicast
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IGMP snooping feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure igmp snooping vlan ports set join-limit

```
configure igmp snooping {vlan} vlanname ports portlist set join-limit  
                        {num}
```

Description

Configures VLAN ports to support a maximum number of IGMP joins.

Syntax Description

<i>vlanname</i>	Specifies a VLAN name.
<i>portlist</i>	Specifies one or more ports or slots and ports.
<i>num</i>	Specifies the maximum number of joins permitted on the ports. The range is 1 to 500.

Default

No limit.

Usage Guidelines

None.

Example

The following example configures port 2:1 in the Default VLAN to support a maximum of 100 IGMP joins:

```
configure igmp snooping "Default" ports 2:1 set join-limit 100
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IGMP snooping feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure igmp ssm-map add

```
configure igmp ssm-map add group_ip [prefix | mask] [source_ip |  
    src_domain_name] {vr vr-name}
```

Description

Configures an [IGMP](#) SSM mapping.

Syntax Description

<i>group_ip</i>	Specifies the multicast IP address for the group mapping.
<i>prefix</i>	Specifies a prefix length for the multicast group IP address. The range is 4 to 32.
<i>mask</i>	Specifies the network mask for the group multicast IP address.
<i>source_ip</i>	The IP address for a multicast group source.
<i>src_domain_name</i>	The source domain name for the multicast group source.
<i>vr-name</i>	Specifies a virtual router name. If the VR name is omitted, the switch uses the VR specified by the current CLI VR context.

Default

N/A.

Usage Guidelines

IGMP SSM mapping operates only with IPv4.

Example

The following example configures an IGMP-SSM mapping for the range of multicast IP addresses at 232.1.1.0/24 originating from IP host 172.16.8.1:

```
configure igmp ssm-map add 232.1.1.0/24 172.16.8.1
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IPv4 multicast feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure igmp ssm-map delete

```
configure igmp ssm-map delete group_ip [prefix] | mask] [source_ip |  
all] vr vr-name
```

Description

Unconfigures an SSM mapping.

Syntax Description

<i>group_ip</i>	Specifies the multicast IP address for the group mapping.
<i>prefix</i>	Specifies a prefix length for the multicast group IP address. The range is 4 to 32.
<i>mask</i>	Specifies the network mask for the group multicast IP address.
<i>source_ip</i>	The IP address for a multicast group source.
all	Specifies that all sources for the specified group or mask are deleted.
<i>vr-name</i>	Specifies a virtual router name. If the VR name is omitted, the switch uses the VR specified by the current CLI VR context.

Default

N/A.

Usage Guidelines

None.

Example

The following example deletes an *IGMP*-SSM mapping for the range of multicast IP addresses at 232.1.1.0/24 originating from IP host 172.16.8.1:

```
configure igmp ssm-map delete 232.1.1.0/24 172.16.8.1
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IPv4 multicast feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure inline-power classification

```
configure inline-power classification [802.3af | 802.3af-high |  
      802.3pre-at | 802.3at | 802.3pre-bt | 802.3bt-type3 | 802.3bt-type4]  
ports port_list
```

Description

This command configures *PoE* port-level classification power-up mode for Extreme Networks PoE devices that do not support 802.3bt when connecting to switches with 60W/90W PoE ports. ExtremeSwitching platforms support per-port basis configuration.

Syntax Description

classification	Classification power-up mode.
802.3af	IEEE 802.3af 15W mode.
802.3af-high	IEEE 802.3af 30W mode.
802.3pre-at	IEEE 802.3 pre-at 30W mode.
802.3at	IEEE 802.3at 30W mode.
802.3pre-bt	IEEE 802.3 pre-bt 4-pair 60W mode.
802.3bt-type3	IEEE 802.3bt Type 3 60W mode.
802.3bt-type4	IEEE 802.3bt Type 4 90W mode.
ports	Specifies the port.
<i>port_list</i>	Port list separated by a comma or - .

Default

Depends on the maximum classification level supported by the platform. For example, the 5720 Series has a default classification level of 802.3bt-type4, while 5320 Series has a default classification level of 802.3at.

Usage Guidelines

Use the specified power-up classification instead of the default classification based on port type.

Example

The following command displays all classification options:

```
# configure inline-power classification ?
802.3af          IEEE 802.3af 15W mode
802.3af-high     IEEE 802.3af 30W mode
802.3at          IEEE 802.3at 30W mode
802.3bt-type3    IEEE 802.3bt Type 3 60W mode
802.3bt-type4    IEEE 802.3bt Type 4 90W mode
802.3pre-at      IEEE 802.3 pre-at 30W mode
802.3pre-bt      IEEE 802.3 pre-bt 4-pair 60W mode
```

History

This command was first available in ExtremeXOS 31.3.

The **802.3af**, **802.3af-high**, and **802.3pre-at** options were added in ExtremeXOS 31.7.

Platform Availability

PoE

PoE+

- ExtremeSwitching 5420F-8W-16P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420F-24P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420F-16MW-32P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420F-16W-32P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420F-48P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420F-48P-4XL—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420M-16MW-32P-4YE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5320-48P-8XE—ExtremeXOS 31.6 and later.
- ExtremeSwitching 5320-24P-8XE—ExtremeXOS 31.6 and later.
- ExtremeSwitching 5320-16P-4XE—ExtremeXOS 31.7 and later.
- ExtremeSwitching 5320-16P-4XE-DC—ExtremeXOS 31.7 and later.

PoE++

- ExtremeSwitching 5520-24W—ExtremeXOS 31.1 and later.
- ExtremeSwitching 5520-48W—ExtremeXOS 31.1 and later.
- ExtremeSwitching 5520-12MW-36W—ExtremeXOS 31.1 and later.
- ExtremeSwitching 5420F-8W-16P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420F-16MW-32P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420F-16W-32P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420M-24W-4YE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420M-16MW-32P-4YE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420M-48W-4YE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5720-24MW—Switch Engine 32.1 and later.
- ExtremeSwitching 5720-24MXW—Switch Engine 32.1 and later.
- ExtremeSwitching 5720-48MW—Switch Engine 32.1 and later.
- ExtremeSwitching 5720-48MXW—Switch Engine 32.1 and later.

configure inline-power detection ports

```
configure inline-power detection [802.3af-only | legacy-and-802.3af [4-point | 2-point] | bypass] ports port_list
```

Description

This command configures *PoE* device detection mode for Extreme Networks PoE devices and SummitStack. ExtremeSwitching platforms support per-port basis configuration.

Syntax Description

802.3af-only	IEEE 802.3af detection only.
legacy-and-802.3af	Capacitive and IEEE 802.3af detection.
4-point	Selects 4-point detection (default).
2-point	Selects 2-point detection (for extended detection signature range).
bypass	No detection phase.
<i>port_list</i>	Port list separated by a comma or - .

Default

Default is **legacy-and-802.3af** detection.

Usage Guidelines

None.

Example

```
# configure inline-power detection ports 1-2
```

History

This command was first available in ExtremeXOS 16.1.

4-point and **2-point** detection options were added in ExtremeXOS 22.5.

Platform Availability

PoE

PoE+

- ExtremeSwitching 5420F-8W-16P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420F-24P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420F-16MW-32P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420F-16W-32P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420F-48P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420F-48P-4XL—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420M-16MW-32P-4YE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5320-48P-8XE—ExtremeXOS 31.6 and later.
- ExtremeSwitching 5320-24P-8XE—ExtremeXOS 31.6 and later.
- ExtremeSwitching 5320-16P-4XE—ExtremeXOS 31.7 and later.
- ExtremeSwitching 5320-16P-4XE-DC—ExtremeXOS 31.7 and later.

PoE++

- ExtremeSwitching 5520-24W—ExtremeXOS 31.1 and later.
- ExtremeSwitching 5520-48W—ExtremeXOS 31.1 and later.
- ExtremeSwitching 5520-12MW-36W—ExtremeXOS 31.1 and later.
- ExtremeSwitching 5420F-8W-16P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420F-16MW-32P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420F-16W-32P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420M-24W-4YE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420M-16MW-32P-4YE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420M-48W-4YE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5720-24MW—Switch Engine 32.1 and later.
- ExtremeSwitching 5720-24MXW—Switch Engine 32.1 and later.
- ExtremeSwitching 5720-48MW—Switch Engine 32.1 and later.
- ExtremeSwitching 5720-48MXW—Switch Engine 32.1 and later.

configure inline-power disconnect-precedence

```
configure inline-power disconnect-precedence [deny-port | lowest-  
priority]
```

Description

Configures the disconnect precedence priority for the switch when a new PD is detected and the measured inline power for that switch or specified slot is within 19 W of the switch's or slot's [PoE](#) power budget.

Syntax Description

deny-port	Specifies power be denied to PD requesting power, regardless of priority.
lowest-priority	Specifies power be withdrawn from lowest-priority port(s) when next PD requesting power connects.

Default

Deny-port.

Usage Guidelines

You configure this parameter for the switch and for the entire SummitStack; you cannot configure this per slot or per port.

If the power supplied to the PDs on a switch or specified slot exceeds the power that was budgeted for that switch or specified slot, the system disconnects power to one or more ports to prevent power overload.

You configure the switch to either deny power to the next PD that requests power on that switch or slot, regardless of the priority, or to disconnect those PDs on ports with lower priorities until there is enough power for the new PD. If you select this last argument and you did not configure port priorities or if several ports have the same priority, the switch withdraws power (or disconnects) those ports with the highest port number (s). For information about configuring the PoE priority for the ports, see [configure inline-power priority ports](#)

The default value is deny-port. So, if you do not change the default value and the switch's or slot's power is exceeded, the next PD requesting power will not be connected.

When the setting is lowest priority, the switch continues dropping ports with the lowest configured PoE port priorities, or the highest port number in the case of equal PoE port priorities, until there is enough power for the requesting PD.

From ExtremeXOS 30.2 , in ExtremeSwitching X465 series switches, when deny port is configured when ports are given priority, priority overtakes deny port action.

Example

The following command sets the switch to withdraw power from the lowest-priority port(s):

```
configure inline-power disconnect-precedence lowest-priority
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on:

PoE+

- ExtremeSwitching 5420F-8W-16P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420F-24P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420F-16MW-32P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420F-16W-32P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420F-48P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420F-48P-4XL—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420M-16MW-32P-4YE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5320-48P-8XE—ExtremeXOS 31.6 and later.
- ExtremeSwitching 5320-24P-8XE—ExtremeXOS 31.6 and later.
- ExtremeSwitching 5320-16P-4XE—ExtremeXOS 31.7 and later.
- ExtremeSwitching 5320-16P-4XE-DC—ExtremeXOS 31.7 and later.

PoE++

- ExtremeSwitching 5520-24W—ExtremeXOS 31.1 and later.
- ExtremeSwitching 5520-48W—ExtremeXOS 31.1 and later.
- ExtremeSwitching 5520-12MW-36W—ExtremeXOS 31.1 and later.
- ExtremeSwitching 5420F-8W-16P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420F-16MW-32P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420F-16W-32P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420M-24W-4YE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420M-16MW-32P-4YE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420M-48W-4YE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5720-24MW—Switch Engine 32.1 and later.
- ExtremeSwitching 5720-24MXW—Switch Engine 32.1 and later.
- ExtremeSwitching 5720-48MW—Switch Engine 32.1 and later.
- ExtremeSwitching 5720-48MXW—Switch Engine 32.1 and later.

configure inline-power label ports

```
configure inline-power label string ports port_list
```

Description

Lets you create your own label for a specified *PoE* port or group of PoE ports.

Syntax Description

<i>string</i>	Specifies a name up to 15 characters in length to identify the specified power port(s).
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

No label.

Usage Guidelines

Use the `show inline-power configuration ports` command, as shown in the following example, to display inline power configuration information, including the label (if any) for each port:

```
show inline-power configuration port 3:1-10
```

Following is sample output from this command on a SummitStack:

Port	Config	Operator Limit	Priority	Label
3:1	Enabled	16000 mW	Low	finance
3:2	Enabled	15000 mW	Low	finance
3:3	Enabled	15000 mW	Low	
3:4	Enabled	15000 mW	Low	
3:5	Enabled	15000 mW	Low	
3:6	Enabled	15000 mW	Low	marketing
3:7	Enabled	15000 mW	Low	marketing
3:8	Enabled	15000 mW	Low	marketing
3:9	Enabled	15000 mW	Low	
3:10	Enabled	15000 mW	Low	

Example

The following command assigns the name “alpha-test_1” to port 1 on slot 4:

```
config inline-power label alpha-test_1 ports 4:1
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on:

PoE+

- ExtremeSwitching 5420F-8W-16P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420F-24P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420F-16MW-32P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420F-16W-32P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420F-48P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420F-48P-4XL—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420M-16MW-32P-4YE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5320-48P-8XE—ExtremeXOS 31.6 and later.
- ExtremeSwitching 5320-24P-8XE—ExtremeXOS 31.6 and later.
- ExtremeSwitching 5320-16P-4XE—ExtremeXOS 31.7 and later.
- ExtremeSwitching 5320-16P-4XE-DC—ExtremeXOS 31.7 and later.

PoE++

- ExtremeSwitching 5520-24W—ExtremeXOS 31.1 and later.
- ExtremeSwitching 5520-48W—ExtremeXOS 31.1 and later.
- ExtremeSwitching 5520-12MW-36W—ExtremeXOS 31.1 and later.
- ExtremeSwitching 5420F-8W-16P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420F-16MW-32P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420F-16W-32P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420M-24W-4YE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420M-16MW-32P-4YE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420M-48W-4YE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5720-24MW—Switch Engine 32.1 and later.
- ExtremeSwitching 5720-24MXW—Switch Engine 32.1 and later.
- ExtremeSwitching 5720-48MW—Switch Engine 32.1 and later.
- ExtremeSwitching 5720-48MXW—Switch Engine 32.1 and later.

configure inline-power operator-limit ports

```
configure inline-power operator-limit [milliwatts | class-based] ports
    [all | port_list]
```

Description

Sets the power limit allowed for PDs connected to the specified ports.

Syntax Description

<i>milliwatts</i>	An integer specifying the maximum allowed power in milliwatts. Default is based on the port's class type (for example, 30W, 60W, 90W).
class-based	Specifies that the power limit is managed by the powered device's class.
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

PoE—15,400 mW.

PoE+—30,000 mW.

PoE++ Type 3—60,000 mW.

PoE++ Type 4—90,000 mW.

Usage Guidelines

This command sets the power limit that a PD can draw on the specified ports. For PoE, the range is 3,000 to 16800mW and the default value is 15,400 mW. For PoE+, the range is 3,000 to 32,000 mW and the default value is 30,000 mW. For PoE++ Type 3, the range is 3,000 mW to 64,000 mW and the default value is 60,000 mW. For PoE++ Type 4, the range is 3,000 mW to 98,000 mW and the default value is 90,000 mW.

If the measured power for a specified port exceeds the port's operator limit, the power is withdrawn from that port and the port moves into a fault state.

If you try to set an operator-limit outside the accepted range, the system returns the following error message:

```
Error: Invalid operator-limit value. Must be in the range of 3000-90000 mW for PoE
802.3bt port
```

With a **class-based** operator-limit type, if a class4 PD is attached, then the operator-limit set will be ignored and the switch will not deliver more than 30W to the PD.

When **class-based** operator-limit is configured on a port, "Max Allowed Power" will show the maximum power for the class of the attached PD. Without the **class-based** option, "Max Allowed Power" will show the configured or default value of operator-limit.

Example

The following command sets the limit for legacy PDs on ports 3–6 of slot 5 on a SummitStack to 10000 mW:

```
configure inline-power operator-limit 10000 ports 5:3-5:6
```

History

This command was first available in ExtremeXOS 11.1.

The **class-based** option was added in ExtremeXOS 32.3.

Platform Availability

This command is available on:

PoE+

- ExtremeSwitching 5420F-8W-16P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420F-24P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420F-16MW-32P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420F-16W-32P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420F-48P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420F-48P-4XL—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420M-16MW-32P-4YE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5320-48P-8XE—ExtremeXOS 31.6 and later.
- ExtremeSwitching 5320-24P-8XE—ExtremeXOS 31.6 and later.
- ExtremeSwitching 5320-16P-4XE—ExtremeXOS 31.7 and later.
- ExtremeSwitching 5320-16P-4XE-DC—ExtremeXOS 31.7 and later.

PoE++

- ExtremeSwitching 5520-24W—ExtremeXOS 31.1 and later.
- ExtremeSwitching 5520-48W—ExtremeXOS 31.1 and later.
- ExtremeSwitching 5520-12MW-36W—ExtremeXOS 31.1 and later.
- ExtremeSwitching 5420F-8W-16P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420F-16MW-32P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420F-16W-32P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420M-24W-4YE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420M-16MW-32P-4YE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420M-48W-4YE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5720-24MW—Switch Engine 32.1 and later.
- ExtremeSwitching 5720-24MXW—Switch Engine 32.1 and later.
- ExtremeSwitching 5720-48MW—Switch Engine 32.1 and later.
- ExtremeSwitching 5720-48MXW—Switch Engine 32.1 and later.

configure inline-power priority ports

```
configure inline-power priority [critical | high | low] ports port_list
```

Description

Sets the *PoE* priority on the specified ports.

Syntax Description

critical high low	Sets the PoE priority for the specified ports.
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

Low.

Usage Guidelines

The system allocates power to those ports with the highest priorities first. This command can also be used in conjunction with the `configure inline-power disconnect-precedence` command. If you configure the disconnect precedence as lowest priority, then newly detected PDs will be powered if that port has higher priority than the existing powered ports.

If there are multiple ports at the same priority level (either configured or by default) and one of the ports must have power withdrawn because of excessive power demands, those ports with the lower port number are powered first. The higher port numbers have power withdrawn first in the case of equal PoE port priorities.

Example

The following command assigns a critical PoE priority on ports 4 – 6 on slot 3 on a SummitSwitch:

```
configure inline-power priority critical ports 3:4-3:6
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on:

PoE+

- ExtremeSwitching 5420F-8W-16P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420F-24P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420F-16MW-32P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420F-16W-32P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420F-48P-4XE—ExtremeXOS 31.3 and later.

- ExtremeSwitching 5420F-48P-4XL—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420M-16MW-32P-4YE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5320-48P-8XE—ExtremeXOS 31.6 and later.
- ExtremeSwitching 5320-24P-8XE—ExtremeXOS 31.6 and later.
- ExtremeSwitching 5320-16P-4XE—ExtremeXOS 31.7 and later.
- ExtremeSwitching 5320-16P-4XE-DC—ExtremeXOS 31.7 and later.

PoE++

- ExtremeSwitching 5520-24W—ExtremeXOS 31.1 and later.
- ExtremeSwitching 5520-48W—ExtremeXOS 31.1 and later.
- ExtremeSwitching 5520-12MW-36W—ExtremeXOS 31.1 and later.
- ExtremeSwitching 5420F-8W-16P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420F-16MW-32P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420F-16W-32P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420M-24W-4YE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420M-16MW-32P-4YE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420M-48W-4YE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5720-24MW—Switch Engine 32.1 and later.
- ExtremeSwitching 5720-24MXW—Switch Engine 32.1 and later.
- ExtremeSwitching 5720-48MW—Switch Engine 32.1 and later.
- ExtremeSwitching 5720-48MXW—Switch Engine 32.1 and later.

configure inline-power usage-threshold

```
configure inline-power usage-threshold threshold
```

Description

Sets the inline power usage [SNMP](#) event threshold.

Syntax Description

<i>threshold</i>	Specifies the percentage of budgeted power used on any PoE switch that causes the system to send an SNMP event and create a log message. The range 1 to 99; the default value is 70.
------------------	--

Default

70.

Usage Guidelines

This command sets the threshold for generating an SNMP event and an EMS message. On a SummitStack, this threshold is when the measured power for a PoE module compared to the budgeted power for that slot exceeds a certain value. On stand-alone switches, this threshold applies to the total power available to the entire switch. The configured threshold value initiates the event and message once that percentage of the budgeted power is being used.

The system generates an additional SNMP event and EMS message once the power usage falls below the threshold again; once the condition clears.

Example

The following command sets the inline power usage alarm threshold at 75%:

```
configure inline-power usage-threshold 75
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on:

PoE+

- ExtremeSwitching 5420F-8W-16P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420F-24P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420F-16MW-32P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420F-16W-32P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420F-48P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420F-48P-4XL—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420M-16MW-32P-4YE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5320-48P-8XE—ExtremeXOS 31.6 and later.
- ExtremeSwitching 5320-24P-8XE—ExtremeXOS 31.6 and later.
- ExtremeSwitching 5320-16P-4XE—ExtremeXOS 31.7 and later.
- ExtremeSwitching 5320-16P-4XE-DC—ExtremeXOS 31.7 and later.

PoE++

- ExtremeSwitching 5520-24W—ExtremeXOS 31.1 and later.
- ExtremeSwitching 5520-48W—ExtremeXOS 31.1 and later.
- ExtremeSwitching 5520-12MW-36W—ExtremeXOS 31.1 and later.
- ExtremeSwitching 5420F-8W-16P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420F-16MW-32P-4XE—ExtremeXOS 31.3 and later.

- ExtremeSwitching 5420F-16W-32P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420M-24W-4YE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420M-16MW-32P-4YE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420M-48W-4YE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5720-24MW—Switch Engine 32.1 and later.
- ExtremeSwitching 5720-24MXW—Switch Engine 32.1 and later.
- ExtremeSwitching 5720-48MW—Switch Engine 32.1 and later.
- ExtremeSwitching 5720-48MXW—Switch Engine 32.1 and later.

configure ip anycast mac

```
configure ip anycast mac [none | mac]
```

Description

Specifies the anycast gateway MAC address that is used by VLANs that enable IP anycast.

Syntax Description

ip	Layer 3 Internet Protocol.
anycast	Configures IP anycast.
mac	Configures the IP anycast MAC address for VLANs with IP anycast enabled.
none	Unconfigures the previously configured MAC address.
<i>mac</i>	Specifies the MAC address.

Default

N/A.

Usage Guidelines

Use this command to specify the same IP address and MAC address on all edge technology devices. This allows seamless IP mobility in the network for edge devices. Using the **none** option unconfigures the previously set MAC address.

Example

The following example specifies the MAC address as 00:00:AB:BA:BA:BA:

```
# configure ip anycast mac 00:00:AB:BA:BA:BA
```

History

This command was first available in ExtremeXOS 30.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ip nat add vlan

```
configure ip nat add {vlan} vlan_name direction [ingress | egress | both]
```

Description

Adds VLANs to the Network Address Translation (NAT).

Syntax Description

ip	Specifies Internet Protocol (IP).
nat	Specifies enabling NAT.
add	Specifies adding a VLAN.
vlan	Specifies adding a VLAN.
<i>vlan_name</i>	Specifies the VLAN to add.
direction	Configures directions for NAT translation on this VLAN.
ingress	Configures VLAN as an ingress VLAN for NAT. Dynamic port translation rules are created for flows ingressing on this VLAN.
egress	Configures VLAN as an egress VLAN for NAT. Dynamic port translation rules are not created for flows ingressing on this VLAN.
both	Configures VLAN for NAT in both directions. Dynamic port translation rules are created for flows ingressing on this VLAN.

Default

N/A.

Usage Guidelines

For NAT Port Translation (NAPT), dynamic NAPT rules are created for traffic coming on VLANs whose direction is configured as “ingress” or “both”.

To delete a VLAN, run the command `configure ip nat delete {vlan} vlan_name`.

Example

The following example adds the VLAN "vlan1" to NAT as an ingress VLAN:

```
# configure ip nat add vlan vlan1 direction ingress
```

History

This command was first available in ExtremeXOS 31.2.

Platform Availability

This command is available on the switches that support the IP NAT feature. For information about which switches support this and other features, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure ip nat aging-time

```
configure ip nat aging-time [minutes | none]
```

Description

Configures the time after which dynamic IP Network Address Translation (NAT) entries that are not active are removed.

Syntax Description

ip	Specifies Internet Protocol (IP).
nat	Specifies NAT.
aging-time	Configures the time after which dynamic IP Network Address Translation (NAT) entries that are not active are removed.
<i>minutes</i>	Specifies the aging time in minutes. The default is 20 minutes. The range is 10 to 3,000 minutes.
none	Specifies not perform aging of NAT entries.

Default

By default, the aging time is 20 minutes.

Usage Guidelines

Example

The following example configures the NAT aging time to 45 minutes.

```
# configure ip nat aging-time 45
```

History

This command was first available in ExtremeXOS 31.2.

Platform Availability

This command is available on the switches that support the IP NAT feature. For information about which switches support this and other features, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure ip nat delete vlan

```
configure ip nat delete {vlan} vlan_name
```

Description

Deletes VLANs from the Network Address Translation (NAT).

Syntax Description

ip	Specifies Internet Protocol (IP).
nat	Specifies enabling NAT.
delete	Specifies deleting a VLAN.
vlan	Specifies deleting a VLAN.
<i>vlan_name</i>	Specifies the VLAN to delete.

Default

N/A.

Usage Guidelines

To add a VLAN, run the command `configure ip nat add {vlan} vlan_name direction [ingress | egress | both]`.

Example

The following example deletes the VLAN "vlan1" from NAT:

```
# configure ip nat delete vlan vlan1
```

History

This command was first available in ExtremeXOS 31.2.

Platform Availability

This command is available on the switches that support the IP NAT feature. For information about which switches support this and other features, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure ip nat rule destination

```
configure ip nat rule rule_name destination [[dst_ip_addr new-destination new_dst_ip_addr {{vr} vr_name}] | none]
```

Description

Configures the IP address that is to be translated and the IP address to which this is to be translated for destination Network Address Port (NAT) rules.

Syntax Description

ip	Specifies Internet Protocol (IP).
nat	Specifies NAT.
rule	Specifies NAT rule.
<i>rule_name</i>	Specifies the NAT rule name.
destination	Specifies modifying the destination IP address.
<i>dst_ip_addr</i>	Specifies the destination IP address that needs to be translated.
new-destination	Specifies a destination IP address after translation.
<i>new_dst_ip_addr</i>	Specifies the destination IP address after translation.
vr	Specifies a virtual router (VR) on which the packet is sent out after translation.
<i>vr_name</i>	Specifies an existing VR/VRF name on which the packet is sent out on after translation. If not specified, the VR context from where this command is executed is used.
none	Deletes the IP address configuration.

Default

If no VR is specified, the current VR context is used.

Usage Guidelines

This command configures the destination NAT rule address information. To configure the source information, run the command `configure ip nat rule rule_name source [[src_ip_addr src_mask | src_ipNetmask] {{source-vr} src_vr_name} new-source new_src_ip_addr] | none`.

Example

The following example configures the rule named "ipOnlyRule" to translate the destination the IP address from 121.144.169.196 to 10.20.30.40. The egress VLAN is present in the VR "VR-user-out":

```
# configure ip nat rule ipOnlyRule destination 121.144.169.196 new-destination 10.20.30.40 vr VR-user-out
```

History

This command was first available in ExtremeXOS 31.2.

Platform Availability

This command is available on the switches that support the IP NAT feature. For information about which switches support this and other features, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure ip nat rule destination protocol

```
configure ip nat rule rule_name destination protocol [[[tcp | udp |
protocol_num] port port_num new-port new_port_num] | none]
```

Description

Configures the protocol that is to be translated and the protocol to which this is to be translated for destination Network Address Translation (NAT) rules.

Syntax Description

ip	Specifies Internet Protocol (IP).
nat	Specifies NAT.
rule	Specifies configuring NAT rules.
<i>rule_name</i>	Specifies the NAT rule name.
destination	Specifies modifying the destination L4 port number.
protocol	Selects the translates IP protocol.
tcp	Specifies Transmission Control Protocol (TCP).
udp	Specifies User Datagram Protocol (UDP).
<i>protocol_num</i>	Specifies the IP protocol number. The range is 1 to 255.
port	Specifies modifying an L4 port.
<i>port_num</i>	Specifies the L4 port number that needs to be translated.
new-port	Specifies the L4 port after translation.
<i>new_port_num</i>	Specifies the L4 port number after translation.
none	Deletes the protocol configuration.

Default

N/A.

Usage Guidelines

Example

The following example specifies for rule "rule1" that the destination protocol is TCP and the port is translated from port 1 to port 2:

```
# configure ip nat rule rule1 destination protocol tcp port 1 new-port 2
```

History

This command was first available in ExtremeXOS 31.2.

Platform Availability

This command is available on the switches that support the IP NAT feature. For information about which switches support this and other features, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure ip nat rule egress

```
configure ip nat rule rule_name egress {vlan} vlan_name
```

Description

Configures the egress VLAN on which the translated traffic is sent out on.

Syntax Description

ip	Specifies Internet Protocol (IP).
nat	Specifies NAT.
rule	Specifies configuring NAT rules.
<i>rule_name</i>	Specifies the NAT rule name.
egress	Specifies configuring the NAT egress VLAN.
vlan	Specifies the NAT egress VLAN.
<i>vlan_name</i>	Specifies the NAT egress VLAN name.

Default

N/A.

Usage Guidelines

To create an IP NAT rule, run the command `create ip nat rule rule_name type [source-nat | napt | destination-napt]`.

Example

The following example configures the VLAN "VLAN1" as the egress VLAN for IP NAT rule "rule2":

```
# configure ip nat rule rule2 egress vlan vlan1
```

History

This command was first available in ExtremeXOS 31.2.

Platform Availability

This command is available on the switches that support the IP NAT feature. For information about which switches support this and other features, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure ip nat rule monitor

```
configure ip nat rule rule_name monitor [on | off]
```

Description

Enables or disables statistics monitoring (byte/packet counters) on a Network Address Translation (NAT) rule in the outward direction and counts Tx counters.

Syntax Description

ip	Specifies Internet Protocol (IP).
nat	Specifies NAT.
rule	Specifies a NAT rule.
<i>rule_name</i>	Specifies the name of the NAT rule to enable monitoring.
monitor	Specifies enabling monitoring.
on	Enables monitoring.
off	Disables monitoring.

Default

N/A.

Usage Guidelines

This command enables/disables statistics monitoring (byte/packet counters) on a NAT rule in the outward direction and counts Tx counters only. There is no provision for Rx counters.

Example

The following example enables monitoring on the NAT rule "rule1":

```
# configure ip nat rule rule1 monitor on
```

The following example disables monitoring on the NAT rule "rule1":

```
# configure ip nat rule rule1 monitor off
```

History

This command was first available in ExtremeXOS 31.2.

Platform Availability

This command is available on the switches that support the IP NAT feature. For information about which switches support this and other features, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure ip nat rule name

```
configure ip nat rule rule_name name new_rule_name
```

Description

Changes the name of an IP Network Address Translation (NAT) rule.

Syntax Description

ip	Specifies Internet Protocol (IP).
nat	Specifies NAT.
rule	Specifies configuring NAT rules.
<i>rule_name</i>	Specifies the current NAT rule name.
name	Specifies providing a new NAT rule name.
<i>new_rule_name</i>	Specifies the new NAT rule name.

Default

N/A.

Usage Guidelines

Example

The following example changes the name of the IP NAT rule "old_rule_name" to "new_rule_name":

```
# configure ip nat rule old_rule_name name new_rule_name
```

History

This command was first available in ExtremeXOS 31.2.

Platform Availability

This command is available on the switches that support the IP NAT feature. For information about which switches support this and other features, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure ip nat rule source

```
configure ip nat rule rule_name source [[src_ip_addr src_mask |
src_ipNetmask ] [{source-vr} src_vr_name} new-source new_src_ip_addr]
| none]
```

Description

Configures the IP address that is to be translated and the IP address to which this is to be translated for source Network Address Translation (NAT) rules.

Syntax Description

ip	Specifies Internet Protocol (IP).
nat	Specifies NAT.
rule	Specifies NAT rule.
<i>rule_name</i>	Specifies the NAT rule name.
source	Specifies modifying the source IP address.
<i>src_ip_addr</i>	Specifies the source IP address that needs to be translated.
<i>src_mask</i>	Specifies source IP address mask.
<i>src_ipNetmask</i>	Specifies source IP address/netmask.
source-vr	Specifies a virtual router (VR) on which the packet to be translated arrives.
<i>src_vr_name</i>	Specifies an existing VR/VRF name on which the packet to be translated arrives. If not specified, the VR context from where this command is executed is used.
new-source	Specifies a source IP address after translation.
<i>new_src_ip_addr</i>	Specifies the source IP address after translation.
none	Deletes the IP address configuration.

Default

If no VR is specified, the current VR context is used.

Usage Guidelines

This command configures the source NAT rule address information. To configure the destination information, run the command `configure ip nat rule rule_name destination [[dst_ip_addr new-destination new_dst_ip_addr {{vr} vr_name}] | none]` .

The incoming and outgoing VLANs on the NAT router can be on different VRs. The VR on which the packet arrives is configured in this command. The packet after translation is sent out on the configured egress VLAN (configure `ip nat add {vlan} vlan_name direction [ingress | egress | both]`).

For a source-NAT rule (where either source or destination IP address is only translated), the source mask has to be specified as /32 or 255.255.255.255.

Example

The following example configures the rule named "ipOnlyRule" to translate the source the IP address from 10.20.30.40 to 121.144.169.196. The ingress VLAN is present in the VR "VR-user-in":

```
# configure ip nat rule ipOnlyRule source 10.20.30.40/32 source-vr VR-user-in new-source 121.144.169.196
```

History

This command was first available in ExtremeXOS 31.2.

Platform Availability

This command is available on the switches that support the IP NAT feature. For information about which switches support this and other features, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure iparp add proxy

```
configure iparp add proxy [ipNetmask | ip_addr {mask}] {vr vr_name} {mac | vrrp} {always}
```

Description

Configures the switch to respond to ARP requests on behalf of devices that are incapable of doing so.

Syntax Description

<i>ipNetmask</i>	Specifies an IP address/mask length.
<i>ip_addr</i>	Specifies an IP address.

<i>mask</i>	Specifies a subnet mask.
<i>vr_name</i>	Specifies a VR name.
<i>mac</i>	Specifies a MAC address to use in the ARP reply.
vrrp	Specifies a MAC address to use in the ARP reply. For VLANs running <i>VRRP</i> , the switch replies with the VRRP virtual MAC. For non-VRRP VLANs, the switch replies with the switch MAC.
always	Specifies that the switch responds regardless of the <i>VLAN</i> that the request arrives from.

Default

If you do not specify a VR or VRF, the current VR context is used.

Usage Guidelines

When *mask* is not specified, an address with the mask 255.255.255.255 is assumed. When neither *mac* nor **vrrp** is 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 source IP address is not on the same subnet as the target address (unless the always flag is set).

After all the proxy ARP conditions have been met, the switch formulates an ARP response using the configured MAC address in the packet.

The default maximum number of proxy entries is 256, but can be increased to 4096 by using the following command:

```
configure iparp max_proxy_entries max_proxy_entries
```

Example

The following example configures the switch to answer ARP requests for all devices with the address range of 100.101.45.1 to 100.101.45.255:

```
configure iparp add proxy 100.101.45.0/24
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure iparp add

```
configure iparp add ip_addr {vr vr_name} mac
```

Description

Adds a permanent entry to the ARP table. You must specify the IP address and MAC address of the entry.

Syntax Description

<i>ip_addr</i>	Specifies an IP address.
<i>vr_name</i>	Specifies a VR name.
<i>mac</i>	Specifies a MAC address.

Default

If you do not specify a VR or VRF, the current VR context is used.

Usage Guidelines

None.

Example

The following example adds a permanent IP ARP entry to the switch for IP address 10.1.2.5:

```
configure iparp add 10.1.2.5 00:11:22:33:44:55
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure iparp delete proxy

```
configure iparp delete proxy [[ipNetmask | ip_addr {mask}] {vr vr_name}  
| all]
```

Description

Deletes one or all proxy ARP entries.

Syntax Description

<i>ipNetmask</i>	Specifies an IP address/mask length.
<i>ip_addr</i>	Specifies an IP address.
<i>mask</i>	Specifies a subnet mask.
<i>vr_name</i>	Specifies a VR name.
all	Specifies all ARP entries.

Default

If you do not specify a VR or VRF, the current VR context is used.

Usage Guidelines

When the mask is not specified, the software assumes a host address (that is, a 32-bit mask).

Example

The following command deletes the IP ARP proxy entry 100.101.45.0/24:

```
configure iparp delete proxy 100.101.45.0/24
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure iparp delete

```
configure iparp delete ip_addr {vr vr_name}
```

Description

Deletes an entry from the ARP table.

Syntax Description

<i>ip_addr</i>	Specifies an IP address.
<i>vr_name</i>	Specifies a VR name.

Default

If you do not specify a VR or VRF, the current VR context is used.

Usage Guidelines

Removes any IP ARP entry (dynamic or permanent) from the table. You must specify the IP address of the entry to delete the entry.

Example

The following command deletes an IP address entry from the ARP table:

```
configure iparp delete 10.1.2.5
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ip-arp fast-convergence

```
configure ip-arp fast-convergence [on | off]
```

Description

This command improves IP convergence for IP traffic.

Syntax Description

on	Fast-convergence on.
off	Fast-convergence off (default).

Default

Off.

Usage Guidelines

Use this command for quick recovery when running IP traffic over an *EAPS* ring.

Example

The following example shows output from the configure ip-arp fast-convergence on command:

```
# show iparp
VR          Destination      Mac           Age  Static  VLAN          VID  Port
VR-Default  10.109.1.2      00:04:96:52:2b:16    0    NO    box1-box2     950   3
VR-Default  10.109.1.6      00:04:96:52:2a:f2    0    NO    box1-box3     951   1
Dynamic Entries :          2          Static Entries :          0
Pending Entries :          0
```

```

In Request      :      1      In Response      :      1
Out Request     :      1      Out Response     :      1
Failed Requests :      0
Proxy Answered  :      0
Rx Error        :      0      Dup IP Addr      :      0.0.0.0
Rejected Count  :
Rejected Port   :
Rejected I/F    :
Max ARP entries :      8192    Max ARP pending entries :      256
ARP address check: Enabled    ARP refresh      :      Enabled
Timeout        :      20 minutes    ARP Sender-Mac Learning :      Disabled
Locktime       :      1000 milliseconds
Retransmit Time :      1000 milliseconds
Reachable Time  :      900000 milliseconds (Auto)
Fast Convergence :      Off

# show iparp
VR      Destination      Mac      Age  Static  VLAN      VID  Port
VR-Default  10.109.1.2      00:04:96:52:2b:16      1      NO  box1-box2  950  3
VR-Default  10.109.1.6      00:04:96:52:2a:f2      1      NO  box1-box3  951  1
Dynamic Entries :      2      Static Entries :      0
Pending Entries :      0
In Request      :      1      In Response      :      1
Out Request     :      1      Out Response     :      1
Failed Requests :      0
Proxy Answered  :      0
Rx Error        :      0      Dup IP Addr      :      0.0.0.0
Rejected Count  :
Rejected Port   :
Rejected I/F    :
Max ARP entries :      8192    Max ARP pending entries :      256
ARP address check: Enabled    ARP refresh      :      Enabled
Timeout        :      20 minutes    ARP Sender-Mac Learning :      Disabled
Locktime       :      1000 milliseconds
Retransmit Time :      1000 milliseconds
Reachable Time  :      900000 milliseconds (Auto)
Fast Convergence :      On

```

History

This command was first available in ExtremeXOS 15.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure iparp locktime

```
configure iparp {vr vr_name}{locktime locktime}
```

Description

Sets the time before a new entry can replace an old entry in the Address Resolution Protocol (ARP) table.

Syntax Description

vr	Specifies setting a VR or VRF.
<i>vr_name</i>	Specifies the name of the VR or VRF.
locktime	Specifies setting a time before a new entry can replace an old entry.
<i>locktime</i>	Sets locktime value (range 0–30,000 milliseconds). Default 1,000 milliseconds.

Default

The default locktime is 1,000 milliseconds.

Example

The following example sets the locktime to 5,000 milliseconds:

```
configure iparp locktime 5000
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure iparp max_entries

```
configure iparp max_entries max_entries
```

Description

Configures the maximum allowed IP ARP entries.

Syntax Description

<i>max_entries</i>	Specifies the maximum number of IP ARP entries. The range is 1 to x, where x is the number listed for the appropriate platform in table below.
--------------------	--

Default

The default value is 12,288, which is the combined value for all VRs, since VR-based maximum entries is not supported starting with ExtremeXOS 30.1.

Usage Guidelines

The maximum IP ARP entries include dynamic, static, and incomplete IP ARP entries. The range for the `max_entries` parameter is 1 to x, where x is the number listed for the appropriate platform in [the following table](#).

Table 10: Maximum IP ARP Entries for each Platform

Maximum Entries	Distributed IP ARP Feature Configuration	
	Off (Default)	On
157,696	N/A	N/A

Starting with ExtremeXOS 30.1, the maximum configurable limit for IP ARP maximum entries is 157,696 for all platforms. A message appears if the configured value exceeds the theoretical hardware maximum limit depending on the platform.

Example

The following example sets the maximum IP ARP entries to 2000 entries:

```
configure iparp max_entries 2000
```

History

This command was first available in ExtremeXOS 10.1.

Support for up to 32,768 ARP entries was first available in ExtremeXOS 12.4.

Per virtual router capability deprecated and the maximum configurable limit changed to 157,696 in ExtremeXOS 30.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure iparp max_pending_entries

```
configure iparp max_pending_entries max_pending_entries
```

Description

Configures the maximum allowed incomplete IP ARP entries.

Syntax Description

<i>max_pending_entries</i>	Specifies a number of maximum IP ARP entries.
----------------------------	---

Default

256.

Usage Guidelines

Range: 1–4,096.

Example

The following example sets the maximum pending IP ARP entries to 500 entries:

```
configure iparp max_pending_entries 500
```

History

This command was first available in ExtremeXOS 10.1.

Per virtual router capability deprecated in ExtremeXOS 30.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure iparp max_proxy_entries

```
configure iparp max_proxy_entries max_proxy_entries
```

Description

Configures the maximum allowed IP ARP proxy entries.

Syntax Description

<i>max_proxy_entries</i>	Specifies maximum number of IP ARP proxy entries.
--------------------------	---

Default

256.

Usage Guidelines

Range: 0–4,096.

Example

The following example sets the maximum IP ARP proxy entries to 500 entries:

```
configure iparp max_proxy_entries 500
```

History

This command was first available in ExtremeXOS 11.0.

Per virtual router capability removed in ExtremeXOS 30.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure iparp proxy reachable | entry-required

```
configure iparp proxy [vlan all | {vlan} vlan_name] [reachable | entry-  
required]
```

Description

Configures whether the switch replies to ARP requests on the specified VLAN by proxy ARP if the route to the IP address is reachable, or only if proxy ARP entries have been created.

Syntax Description

vlan	Selects VLAN(s) for the ARP requests.
<i>vlan_name</i>	Specifies VLAN name for the ARP requests.
all	Specifies all VLANs for the ARP requests.
reachable	Specifies that the switch replies to ARP requests on the specified VLAN(s) by proxy ARP if the route to IP address is reachable. Configuration of proxy ARP entries is not required.
entry-required	Specifies that the switch replies to ARP requests on the specified VLAN(s) by proxy ARP only if proxy ARP entries have been created. (Default)

Default

The default behavior is for the switch to reply to ARP requests on the specified VLAN(s) by proxy ARP only if proxy ARP entries have been created.

Usage Guidelines

If an ARP request is received by the switch, it checks the ExtremeXOS proxy ARP table (user adds the entries through the CLI). If it is present, an ARP reply is sent. If not present, it searches for the entry in the kernel route table. If this IP address is reachable, then the ARP reply is sent.

The following table summarizes the this command's behavior:

	reachable	entry-required (or command not configured)
Entry present in proxy ARP Table (static entry added through command configure iparp add on page 624)	Reply to ARP request	Reply to ARP request
Static entry not present in proxy ARP table	Reply to ARP request if route is reachable	No reply to ARP request
No static entry, but route reachable.	Reply to ARP request	
No static entry and route is not reachable	No reply to ARP request	
Static entry present and route is reachable	Reply to ARP request	
No static entry and route is not reachable	No reply to ARP request	
Route reachable	Reply to ARP request	
Route not reachable	No reply to ARP request	

Example

The following example configures the switch to reply to ARP requests on all VLANs by proxy ARP if the route to IP address is reachable:

```
configure iparp proxy vlan all reachable
```

History

This command was first available in ExtremeXOS 22.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure iparp reachable-time

```
configure iparp {reachable-time [auto | {reachable_time [seconds | milliseconds] } ] }
```

Description

Sets the value for Address Resolution Protocol (ARP) reachable time

Syntax Description

reachable-time	Specifies setting the ARP reachable time.
<i>reachable_time</i>	Sets the value for the ARP reachable time (range is 1-1,474,515,000 milliseconds or 1-1,474,515 seconds).
auto	Specifies having the ARP reachable time set automatically to 3/4 of the configured ARP timeout (default).
milliseconds	When setting the reachable time value, specifies milliseconds as the time unit (range is 1-1,474,515,000).
seconds	When setting the reachable time value, specifies seconds (range is 1-1,474,515) as the time unit (default).

Default

The default setting is for the reachable time to be set automatically to 3/4 of the configured ARP timeout. If you set the time manually, the default unit of measure for the value is seconds.

Example

The following example sets the reachable time to 500,000 seconds:

```
configure iparp reachable-time 500000 seconds
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure iparp retransmit-time

```
configure iparp {retransmit-time retransmit_time}
```

Description

Sets the value for Address Resolution Protocol (ARP) retransmit time

Syntax Description

retransmit-time	Specifies setting the retransmit time.
<i>retransmit_time</i>	Sets the retransmit time value (range is 1-4,294,967 seconds or 1-4,294,967,295 milliseconds). Default is 1 second.
milliseconds	When setting the retransmit time value, specifies milliseconds as the time unit (range is 1-4,294,967,295).
seconds	When setting the retransmit time value, specifies seconds (range is 1-4,294,967) as the time unit (default).

Default

The default setting for the retransmit time is 1 second. The default unit of measure is seconds.

Example

The following example sets the retransmit time to 500,000 seconds:

```
configure iparp retransmit-time 500000 seconds
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure iparp timeout

```
configure iparp timeout {vr vr_name} minutes
```

Description

Configures the IP ARP timeout period.

Syntax Description

<i>vr_name</i>	Specifies which VR or VRF IP ARP setting to change.
<i>minutes</i>	Specifies a time in minutes.

Default

20 minutes.

If you do not specify a VR or VRF, the current VR context is used.

Usage Guidelines

The range is 0-32,767. A setting of 0 disables timeout.

When the switch learns an ARP entry, it begins the timeout for that entry. When the timer reaches 0, the entry is aged out, unless IP ARP refresh is enabled. If ARP refresh is enabled, the switch sends an ARP request for the address before the timer expires. If the switch receives a response, it resets the timer for that address.

Newly configured ARP timeout values apply only to ARP entries that are learned after the new value is set. Previously learned ARP entries timeout after the previously configured time.

Example

The following command sets the IP ARP timeout period to 10 minutes:

```
configure iparp timeout 10
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ipforwarding originated-packets

```
configure ipforwarding originated-packets [require-ipforwarding | dont-require-ipforwarding]
```

Description

Configures whether IP forwarding must be enabled on a VLAN before transmitting IP packets originated by the switch on that VLAN to a gateway.

Syntax Description

require-ipforwarding	Specifies that IP forwarding must be enabled on a VLAN before IP packets that originate on the switch can be transmitted to a gateway.
dont-require-ipforwarding	Specifies that all IP packets that originate on the switch can be transmitted, regardless of the IP forwarding configuration to the gateway.

Default

dont-require-ipforwarding.

Usage Guidelines

To display the current setting for this command, use the `show ipconfig` command.

Example

The following command configures the switch to transmit switch-originated packets to gateways only on those VLANs for which IP forwarding is enabled:

```
configure ipforwarding originated-packets require-ipforwarding
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms that use the Edge, Advanced Edge, or Core license. For information on the licenses available for each platform, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure ipmcforwarding

```
configure ipmcforwarding to-cpu [auto | off] ports port_list
```

Description

Configure whether IP multicast CPU filters are installed automatically.

Syntax Description

auto	The software will automatically program IP multicast processing based on configuration.
off	IP multicast packets received on this port are always flooded with no CPU processing.
<i>port_list</i>	Specifies on or more ports.

Default

N/A.

Usage Guidelines

IP forwarding and IPMC forwarding must be enabled for the configuration to operate.

Example

The following example configures automatic operation for port 2.1:

```
configure ipmcforwarding to-cpu auto ports 2.1
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IPv4 multicast feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure ipmroute add

```
configure ipmroute add [default | source-net mask-len | source-net mask]  
                        [{protocol} protocol] rpf-address {metric} {vr vr-name}
```

Description

Adds a static multicast route to the multicast routing table.

Syntax Description

default	Specifies default gateway.
<i>source-net</i>	Specifies an IP address/mask length.
<i>mask-len</i>	Mask length for the IP multicast source's subnet. Range is [1-32].
<i>mask</i>	Specifies a subnet mask.
<i>protocol</i>	Unicast routing protocol that is to be used for route learning.
<i>rpf-address</i>	Next hop through which the multicast source can be reached.
<i>metric</i>	Specifies a cost metric.
<i>vr-name</i>	Specifies the virtual router to which the route is added.

Default

The following defaults apply:

- metric—1
- vr-name—VR of the current CLI context
- protocol—none

Usage Guidelines

This command allows you to statically configure where multicast sources are located (even though the unicast routing table has different entries). It allows you to configure a multicast static route in such a way as to have non-congruent topology for Unicast and Multicast topology and traffic.

Example

The following example configures a multicast static route for all multicast sources within network subnet 192.168.0.0/16. Those sources are reachable through the gateway 192.75.0.91.

```
configure ipmroute add 192.168.0.0/16 192.75.0.91
```

The following example configures multicast static route for all sources via a single gateway with a metric of 100:

```
configure ipmroute add 0.0.0.0/0 192.75.0.91 100
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IPv4 multicast feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure ipmroute delete

```
configure ipmroute delete [default | source-net/mask-len | source-net mask] [{protocol} protocol] rpf-address {vr vr-name}
```

Description

Deletes a static multicast address from the multicast routing table.

Syntax Description

default	Specifies default gateway.
<i>source-net</i>	Specifies an IP address/mask length.
<i>mask-len</i>	Mask length for the IP multicast source's subnet. Range is 1-32.
<i>mask</i>	Specifies a subnet mask.
<i>protocol</i>	Unicast routing protocol that is to be used for route learning.
<i>rpf-address</i>	Next hop through which the multicast source can be reached.
<i>vr-name</i>	Specifies the virtual router to which the route is added.

Default

vr-name is the VR of the current CLI context.

Usage Guidelines

This command allows you to delete an existing multicast static route. It allows you to configure congruent topology for unicast and multicast packets and traffic.

Example

The following example deletes a multicast static route:

```
configure ipmroute delete 192.168.0.0/16 192.75.0.91
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IPv4 multicast feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure ip-mtu vlan

```
configure ip-mtu mtu [ {vlan} vlan_name | vlan vlan_list]
```

Description

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

Syntax Description

<i>mtu</i>	Specifies the IP maximum transmission unit (MTU) value. Range is from 1500 to 9194. However, CLI will allow the maximum limit upto 9216 considering port configuration such as tagging which influences L2 Header size. But the values greater than 9194 may lead to packet loss and hence not recommended.
<i>vlan_name</i>	Specifies a VLAN name.
<i>vlan_list</i>	Specifies a VLAN list of IDs.

Default

The default IP MTU size is 1500.

Usage Guidelines

Use this command to enable jumbo frame support or for IP fragmentation with jumbo frames. Jumbo frames are Ethernet frames that are larger than 1522 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 4bytes of CRC included in a jumbo frame configuration. 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.

Example

The following example sets the MTU size to 2000 for VLAN sales:

```
configure ip-mtu 2000 vlan sales
```

History

This command was available in ExtremeXOS 11.0.

The `vlan_list` option was added in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure iproute add blackhole ipv4 default

```
configure iproute add blackhole ipv4 default {multicast | multicast-only  
| unicast | unicast-only} {vr vrname}
```

Description

Adds a default blackhole route to the routing table. All traffic destined for an unknown IP destination is silently dropped, and no *ICMP (Internet Control Message Protocol)* message is generated.

Syntax Description

multicast	Adds the default blackhole route to the multicast routing table.
multicast-only	Adds the default blackhole route to the multicast routing table. This option is provided for backward compatibility with releases prior to ExtremeXOS Release 12.1.
unicast	Adds the default blackhole route to the unicast routing table.
unicast-only	Adds the default blackhole route to the unicast routing table. This option is provided for backward compatibility with releases prior to ExtremeXOS Release 12.1.
<i>vrname</i>	Specifies the VR or VRF to which the route is added.

Default

If you do not specify a VR or VRF, the current VR context is used.

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.

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 example adds a blackhole default route into the routing table:

```
configure iproute add blackhole default
```

History

This command was first available in ExtremeXOS 10.1.

The ipv4 keyword was added in ExtremeXOS 11.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure iproute add blackhole ipv6 default

```
configure iproute add blackhole ipv6 default {vr vr_name} {multicast-only | unicast-only}
```

Description

Adds a default blackhole route to the routing table. All traffic destined for an unknown IPv6 destination is silently dropped.

Syntax Description

<i>vr_name</i>	Specifies the VR or VRF to which the route is added.
multicast-only	Specifies only multicast traffic for the route.
unicast-only	Specifies only unicast traffic for the route.

Default

If you do not specify a VR or VRF, the current VR context is used.

Usage Guidelines

While a default route is for forwarding traffic destined to an unknown IPv6 destination, and a blackhole route is for discarding traffic destined to a specified IPv6 destination, a default blackhole route is for discarding traffic to the unknown IPv6 destination.

Using this command, all traffic with an unknown destination is discarded.

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 IPv6 address for this route is ::.

The packets are silently discarded. In other words, no ICMP message is sent to indicate that the packets are discarded.

Example

The following example adds a blackhole default route into the routing table:

```
configure iproute add blackhole ipv6 default
```

History

This command was first available in ExtremeXOS 10.1.

Support for IPv6 was added in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure iproute add blackhole

```
configure iproute add blackhole {ipv6} [ipv6Netmask] {vr vr_name}  
    {multicast-only | unicast-only}
```

Description

Adds a blackhole address to the routing table. All traffic destined for an unknown IPv6 destination is silently dropped.

Syntax Description

<i>ipv6Netmask</i>	Specifies an IPv6 address/prefix length.
<i>vr_name</i>	Specifies the VR or VRF to which the route is added.
multicast-only	Specifies only multicast traffic for the route.
unicast-only	Specifies only unicast traffic for the route.

Default

If you do not specify a VR or VRF, the current VR context is used.

Usage Guidelines

A blackhole entry directs packets with a matching specified address prefix 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.

The packets are silently discarded. In other words, no *ICMP* message is sent to indicate that the packets are discarded.

Example

The following example causes packets with a destination address of 2001:db8::3452 to be silently discarded:

```
configure iproute add blackhole 2001:db8::3452/128
```

History

This command was first available in ExtremeXOS 10.1.

Support for IPv6 was added in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure iproute add default

```
configure iproute add default [gateway | ipv6Gateway | ipv6ScopedGateway]  
    {bfd} {metric} {vr vr_name} {multicast | multicast-only | unicast |  
    unicast-only} {vlan vlan_name}
```

Description

Adds a default gateway to the routing table.

Syntax Description

<i>gateway</i>	Specifies a gateway IPv4 address.
<i>ipv6Gateway</i>	Specifies a <i>VLAN</i> gateway IPv6 address.
<i>ipv6ScopedGateway</i>	Specifies a scoped gateway.

bfd	Enables Bidirectional Forwarding Detection (BFD) protection for the route. Note: You must type this keyword before specifying a VLAN.
metric	Specifies a cost metric. If no metric is specified, the default of 1 is used.
<i>vr_name</i>	Specifies the VR or VRF to which the route is added.

Default

If no metric is specified, the default metric of 1 is used. If you do not specify a VR or VRF, the current VR context 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 IPv6 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 example configures a default route for the switch:

```
configure iproute add default 2001:db8::1234:5678
```

History

This command was first available in ExtremeXOS 10.1.

Support for IPv6 was added in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure iproute add (IPv4)

```
configure iproute add [ipNetmask | ip_addr mask] gateway {bfd}  
                    {metric} {multicast | multicast-only | unicast | unicast-only} {vlan  
                    egress_vlan} {vr vrname}
```

Description

Adds a static route to the specified routing table.

Syntax Description

<i>ipNetmask</i>	Specifies an IP address/mask length.
<i>ip_addr</i>	Specifies an IP address.
<i>mask</i>	Specifies a subnet mask.
<i>gateway</i>	Specifies a gateway IP address.
bfd	Enables Bidirectional Forwarding Detection (BFD) protection for the route. Note: You must type this keyword before specifying a VLAN.
metric	Specifies a cost metric.
multicast	Adds the specified route to the multicast routing table.
multicast-only	Adds the specified route to the multicast routing table. This option is provided for backward compatibility with releases prior to ExtremeXOS release 12.1.
unicast	Adds the specified route to the unicast routing table.
unicast-only	Adds the specified route to the unicast routing table. This option is provided for backward compatibility with releases prior to ExtremeXOS Release 12.1.
vlan	Specifies the egress <u>VLAN</u> name used for an Inter-VR route.
<i>vrname</i>	Specifies the VR or VRF to which the route is added.

Default

If you do not specify a VR or VRF, the current VR context is used.

Usage Guidelines

Use a mask value of 255.255.255.255 to indicate a host entry.

The gateway address must be present on a directly attached subnet, or the following message appears:

```
ERROR: Gateway is not on directly attached subnet
```

The gateway address must be different from loop back address or local addresses, or the following message appears:

```
ERROR: Gateway cannot be local or loop back address
```



Note

Although dynamic unicast routes can be captured in the multicast routing table, unicast static routes cannot be captured in the multicast routing table. To create a static route for the multicast routing table, you must specify the multicast option.

This command can add BFD protection to a link only when the BFD client at each end of the link is enabled (see the [configure iproute add \(IPv4\)](#) command).

Once the BFD session is established, the operational status of the route reflects the operational status of the BFD session.

To remove BFD protection for a static route, enter this command without the BFD keyword.

Beginning in ExtremeXOS 15.6, the egress VLAN name may now be a VLAN belonging to a VR different from the VR of the static route itself. When the VRs differ, Inter-VR routing of hardware and software forwarded packets is performed.

Example

The following example adds a static address to the routing table in the current VR context:

```
configure iproute add 10.1.1.0/24 123.45.67.1
```

In the following example of an Inter-VR routing scenario, VLAN v1 belongs to VR vr1, and VLAN v2 belongs to VR vr2. The final two commands add Inter-VR routes between VR vr1 and VR vr2. The resulting behavior is that IPv4 unicast packets originating in VR vr1, and a destination IP address in subnet 52.0.0.0/8, are forwarded to gateway 20.1.1.2 belonging to VLAN v2 in VR vr2 per the first Inter-VR route. Reverse packets originating in VR vr2 with a destination IP address in subnet 51.0.0.0/8 are forwarded to gateway 10.1.1.2 belonging to VLAN v1 in VR vr1 per the second Inter-VR route. The **vr_name** of the static route command refers to which VR's route table the route is added.

```
create vr "vr1"
create vr "vr2"
create vlan "v1" vr vr1
create vlan "v2" vr vr2
configure vlan v1 tag 10
configure vlan v2 tag 20
configure vlan v1 add ports 1 tagged
configure vlan v2 add ports 2 tagged
configure vlan v1 ipaddress 10.1.1.1/8
configure vlan v2 ipaddress 20.1.1.1/8
enable ipforwarding vlan v1
enable ipforwarding vlan v2

configure iproute add 52.0.0.0/8 20.1.1.2 vlan v2 vr vr1
configure iproute add 51.0.0.0/8 10.1.1.2 vlan v1 vr vr2
```

The Inter-VR routing example above is for packets routed through a gateway to a remote subnet. Inter-VR routing can also be accomplished to/from a host adjacent to the switch, such as hosts in the switch's IPv4 ARP cache, by adding a /32 host route. In the example network above, to have packets from VR1 route to a host/server in VR2 directly on the 20.1.1/8 subnet, such as 20.1.1.66, the following CLI command can be used by specifying 20.1.1.66/32:

```
configure iproute add 20.1.1.66/32 20.1.1.66 vlan v2 vr vr1
```

History

This command was first available in ExtremeXOS 10.1.

Beginning in ExtremeXOS 15.6, the egress VLAN name may now be a VLAN belonging to a VR different from the VR of the static route itself.

Platform Availability

This command is available on all platforms with Layer 3 support.

configure iproute add (IPv6)

```
configure iproute add ipv6Netmask [ipv6Gateway | ipv6ScopedGateway]
    {bfd} {metric} {vr vr_name} {multicast | multicast-only | unicast |
    unicast-only}
```

Description

Adds an IPv6 static route to the routing table.

Syntax Description

<i>ipv6Netmask</i>	Specifies an IPv6 address/prefix length.
<i>ipv6Gateway</i>	Specifies a gateway.
<i>ipv6ScopedGateway</i>	Specifies a scoped gateway.
bfd	Enables Bidirectional Forwarding Detection (BFD) protection for the IPv6 route.
<i>metric</i>	Specifies a cost metric.
<i>vr_name</i>	Specifies the VR or VRF to which the route is added.
multicast	Adds the specified route to the multicast routing table.
multicast-only	Adds the specified route to the multicast routing table. This option is provided for backward compatibility with releases prior to ExtremeXOS release 12.1.
unicast	Adds the specified route to the unicast routing table.
unicast-only	Adds the specified route to the unicast routing table. This option is provided for backward compatibility with releases prior to ExtremeXOS release 12.1.

Default

If you do not specify a VR or VRF, the current VR context is used. If you do not specify a metric, then the default metric of 1 is used.

Usage Guidelines

Use a prefix length of 128 to indicate a host entry.



Note

Although dynamic unicast routes can be captured in the multicast routing table, unicast static routes cannot be captured in the multicast routing table. To create a static route for the multicast routing table, you must specify the multicast option.

Example

The following example adds a static route to the routing table:

```
configure iproute add 2001:db8:0:1111::/64 fe80::1111%default
```

History

This command was first available in ExtremeXOS 10.1.

Support for IPv6 was added in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure iproute add lsp

```
configure iproute add [ipaddress netmask | ipNetmask] lsp lsp_name
    {metric} {multicast | multicast-only | unicast | unicast-only} {vr
    vrname}
```

Description

Assigns a specific IP route to use a named LSP.



Note

To create a static IP route that does not use a specific named LSP as an mpls-next-hop, use the following command: `configure iproute add [ipNetmask | ip_addr mask] gateway {metric} {multicast | multicast-only | unicast | unicast-only} {vr vrname} .`

Syntax Description

<i>ipaddress</i>	Specified an IP address.
<i>netmask</i>	Specifies an IP address/prefix length.
<i>ipNetmask</i>	Specifies an IP address/prefix length.
<i>lsp_name</i>	Specifies a named <i>MPLS</i> LSP to be used to reach the route.
metric	Specifies a cost metric.
multicast	Adds the specified route to the multicast routing table.
multicast-only	Adds the specified route to the multicast routing table. This option is provided for backward compatibility with releases prior to ExtremeXOS Release 12.1.
unicast	Adds the specified route to the unicast routing table.

unicast-only	Adds the specified route to the unicast routing table. This option is provided for backward compatibility with releases prior to ExtremeXOS release 12.1.
<i>vrname</i>	Specifies the virtual router to which the route is added.

Default

N/A.

Usage Guidelines

This command assigns a named LSP to a specific IP route. Once configured, all IP traffic matching the configured route is forwarded over the specified LSP. For an RSVP-TE LSP, the correct label information is only associated with the route if the LSP is active. If the RSVP-TE LSP is disabled or is withdrawn, the label information is removed from the route table and the route entry is marked down. If multiple LSPs are added to a route and *ECMP* is enabled using route-sharing command, only one LSP is used to forward IP traffic.



Note

IP routes can only be assigned to named LSPs in the VR in which MPLS is configured to operate.

Example

The following command adds a static address to the routing table:

```
configure iproute add 10.1.1.0/24 lsp lsp598
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure iproute add (Multicast)

```
configure iproute add [ipNetmask | ip_addr mask] gateway {bfd} {metric}
                    {multicast | multicast-only | unicast | unicast-only} {vr vrname}
```

Description

Adds a static route to the routing table.

Syntax Description

<i>ipNetmask</i>	Specifies an IP address/mask length.
<i>ip_addr</i>	Specifies an IP address.
<i>mask</i>	Specifies a subnet mask.
<i>gateway</i>	Specifies a <i>VLAN</i> gateway.
bfd	Enables Bidirectional Forwarding Detection (BFD) protection for the IPv6 route.
metric	Specifies a cost metric.
<i>vrname</i>	Specifies the virtual router to which the route is added.
multicast	Adds the specified route to the multicast routing table.
multicast-only	Adds the specified route to the multicast routing table. This option is provided for backward compatibility with releases prior to ExtremeXOS 12.1.
unicast	Adds the specified route to the unicast routing table.
unicast-only	Adds the specified route to the unicast routing table. This option is provided for backward compatibility with releases prior to ExtremeXOS 12.1.

Default

If you do not specify a virtual router, the current virtual router context is used.

Usage Guidelines

Use a mask value of 255.255.255.255 to indicate a host entry.



Note

Although dynamic unicast routes can be captured in the multicast routing table, unicast static routes cannot be captured in the multicast routing table. To create a static route for the multicast routing table, you must specify the multicast option.

Example

The following example adds a static address to the multicast routing table:

```
configure iproute add 10.1.1.0/24 123.45.67.1 5 multicast
```

History

This command was first available in ExtremeXOS 10.1.

The **multicast** and **unicast** keywords were first available in ExtremeXOS 12.1. These keywords replace **multicast-only** and **unicast-only**, which remain in the software for backward compatibility.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IPv4 multicast feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure iproute add protection

```
configure iproute add [default | ipv4_or_ipv6_network] gateway  
    {protection [bfd | ping | none]}
```

Description

Configures protection and resiliency on IPv4 and IPv6 static routes.

Syntax Description

default	Default route.
<i>ipv4_or_ipv6_network</i>	IPv4 or IPv6 network address.
<i>gateway</i>	Gateway IP address.
protection	Selects the type of protection on this route (default is none).
bfd	Enables BFD protection on this route.
ping	Enables <i>ICMP</i> ping protection on this route.
none	Disables all protection on this route (default).

Default

No protection is the default.

Usage Guidelines

For static routes configured with protection type ping, static routes are initially down. Static routes become "up" for each configured gateway/device IP when a timely ICMP Echo Reply is received from that IP within the configured ping interval. Static routes transition from up to down when no timely reply is received for the configured number of missed intervals. Severely delayed ICMP Echo Replies are ignored if received after the configured interval time elapses, because a new ICMP Echo Request has already been sent. Static routes with ping protection need not be *ECMP* routes. Thus when a device is unresponsive, a different route with a higher cost or shorter prefix length can route packets elsewhere.

The protection type (BFD, ping, or none) for an existing static route can be changed dynamically without deleting the route. To change the protection type, simply re-add an existing static route with a different protection type.

Example

The following example adds a static route for 100.0.0.0/24 with ping health check monitoring to gateway IP 1.2.3.4.

```
# configure iproute add 100.0.0.0/24 1.2.3.4 protection ping
```

ExtremeXOS initiates ping health check monitoring to the adjacent device with IP address 1.2.3.4. The route for 100.0.0.0/24 is protected, meaning if ping responses are received from 1.2.3.4 in a timely manner, the static route for 100.0.0.0/24 to 1.2.3.4 is “up” in the routing table. If no ping response is received in a timely manner, the route is down.

In an example with ECMP, assuming `enable iproute sharing`:

```
# configure iproute add 100.0.0.0/24 1.2.3.5 protection ping
```

If ping responses are received by both 1.2.3.4 and 1.2.3.5, IP packets destined to subnet 100.0.0.0/24 are Layer-3 load balanced by hardware between 1.2.3.4 and 1.2.3.5. If for example, no ping response is received from 1.2.3.4 in a timely manner, IP packets destined to 100.0.0.0/24 are sent only to 1.2.3.5. Later, upon receiving a ping response from 1.2.3.4, packets are load balanced again.

History

This command was first available in ExtremeXOS 22.1.

Platform Availability

This command is available on all platforms with a Base license or higher as described in the [***Switch Engine 32.3 Feature License Requirements***](#) document.

configure iproute delete

```
configure iproute delete [ipNetmask | ipaddress mask] gateway {multicast
| multicast-only | unicast | unicast-only} {vlan egress vlan} {vr
vrname}
```

Description

Deletes a static address from the routing table.

Syntax Description

<i>ipNetmask</i>	Specifies an IP address/mask length.
<i>ipaddress</i>	Specifies an IP address.
<i>mask</i>	Specifies a subnet mask.
<i>gateway</i>	Specifies a <u>VLAN</u> gateway.
multicast	Specifies a multicast route to delete.
multicast-only	Specifies a multicast route to delete.

unicast	Specifies a unicast route to delete.
unicast-only	Specifies a unicast route to delete.
vlan	Specifies the egress VLAN name used for an Inter-VR route.
<i>vrname</i>	Specifies the virtual router to which the route is deleted.

Default

If you do not specify a virtual router, the current virtual router context is used.

Usage Guidelines

Use a value of 255.255.255.255 or /32 for mask to indicate a host entry.

Example

The following example deletes an address from the multicast routing table:

```
configure iproute delete 10.101.0.0/24 10.101.0.1 multicast
```

History

This command was first available in ExtremeXOS 10.1.

The **multicast** and **unicast** keywords were first available in ExtremeXOS 12.1. These keywords replace **multicast-only** and **unicast-only**, which remain in the software for backward compatibility.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IPv4 multicast feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure iproute delete blackhole

```
configure iproute delete blackhole [ipv6Netmask] {vr vr_name}
```

Description

Deletes a blackhole route from the routing table.

Syntax Description

<i>ipv6Netmask</i>	Specifies an IPv6 address/prefix length.
<i>vr_name</i>	Specifies the VR or VRF from which the route is deleted.

Default

If you do not specify a VR or VRF, the current VR context is used.

Usage Guidelines

A blackhole entry directs packets with a specified 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.

Example

The following example deletes a blackhole route from the routing table for packets with a destination address of 2001:db8::3452, so the packets are no longer discarded:

```
configure iproute delete blackhole 2001:db8::3452/128
```

History

This command was first available in ExtremeXOS 10.1.

Support for IPv6 was added in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure iproute delete blackhole ipv4 default

```
configure iproute delete blackhole ipv4 default {multicast | multicast-only | unicast | unicast-only} {vr vrname}
```

Description

Deletes a default blackhole route from the routing table.

Syntax Description

multicast	Specifies a default blackhole multicast route to delete.
multicast-only	Specifies a default blackhole multicast route to delete. This option is provided for backward compatibility with releases prior to ExtremeXOS Release 12.1.
unicast	Specifies a default blackhole unicast route to delete.

unicast-only	Specifies a default blackhole unicast-only route to delete. This option is provided for backward compatibility with releases prior to ExtremeXOS Release 12.1.
<i>vrname</i>	Specifies a VR or VRF name.

Default

If you do not specify a VR or VRF, the current VR context is used.

Usage Guidelines

None.

Example

The following command deletes a blackhole default route from the routing table:

```
configure iproute delete blackhole default
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure iproute delete blackhole ipv6 default

```
configure iproute delete blackhole ipv6 default {vr vr_name}
```

Description

Deletes a default blackhole route from the routing table.

Syntax Description

<i>vr_name</i>	Specifies the VR or VRF from which the route is deleted.
----------------	--

Default

If you do not specify a VR or VRF, the current VR context is used.

Usage Guidelines

While a default route is for forwarding traffic destined to an unknown IPv6 destination, and a blackhole route is for discarding traffic destined to a specified IPv6 destination, a default blackhole route is for discarding traffic to the unknown IPv6 destination.

Using this command, all traffic with an unknown destination is discarded.

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 IPv6 address for this route is ":::"

Example

The following example deletes a blackhole default route from the routing table:

```
configure iproute delete blackhole default
```

History

This command was first available in ExtremeXOS 10.1.

Support for IPv6 was added in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure iproute delete default

```
configure iproute delete default [ipv6Gateway | ipv6ScopedGateway] {vr  
    vr_name}
```

Description

Deletes a default gateway from the routing table.

Syntax Description

<i>ipv6Gateway</i>	Specifies a <u>VLAN</u> gateway IPv6 address.
<i>ipv6ScopedGateway</i>	Specifies a scoped gateway.
<i>vr_name</i>	Specifies the VR or VRF from which the route is deleted.

Default

If no metric is specified, the default metric of 1 is used. If you do not specify a VR or VRF, the current VR context 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 IPv6 interface.

Example

The following example deletes a default route from the switch:

```
configure iproute delete default 2001:db8::1234:5678
```

History

This command was first available in ExtremeXOS 10.1.

Support for IPv6 was added in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure iproute ipv6 priority

```
configure iproute ipv6 priority [auto-peering | ripng | blackhole | icmp
| host-mobility | static | ospfv3-intra | ospfv3-inter | ospfv3-as-external
| ospfv3-extern1 | ospfv3-extern2 | isis | isis-level-1 | isis-level-2
| isis-level-1-external | isis-level-2-external | ebgp | ibgp] priority {vr vr_name}
```

Description

Changes the priority for all routes from a particular route origin.

Syntax Description

auto-peering	Specifies auto-peering routes.
ripng	Specifies <i>RIPng</i> .
host-mobility	Host-Mobility route.
blackhole	Specifies the blackhole route.
icmp	Specifies <i>ICMP</i> .
static	Specifies static routes.
ospfv3-intra	Specifies <i>OSPFv3</i> Intra routing.
ospfv3-inter	Specifies OSPFv3 Inter routing.
ospfv3-as-external	Specifies OSPFv3 AS External routing.
ospfv3-extern1	Specifies OSPFv3 External 1 routing.

ospfv3-extern2	Specifies OSPFv3 External 2 routing.
isis	Specifies ISIS routing.
isis-level-1	Specifies IS-IS Level 1 routing.
isis-level-2	Specifies IS-IS Level 2 routing.
isis-level-1-external	Specifies IS-IS Level 1 External routing.
isis-level-2-external	Specifies IS-IS Level 2 External routing.
ebgp	Specifies EBGp routes.
ibgp	Specifies IBGP routes.
<i>priority</i>	Specifies a priority number in the range of 11 to 65534.
<i>vr_name</i>	Specifies a VR or VRF name.

Default

The following table lists the relative priorities assigned to routes depending upon the learned source of the route.

Table 11: Route Priorities

Route Origin	Priority
Direct	10
BlackHole	50
Static	1100
HostMobility	1150
ICMP	1200
OSPF3Intra	2200
OSPF3Inter	2300
IS-IS L1	2360
IS-IS L2	2370
RIPg	2400
OSPFv3 ASExt	3100
OSPFv3 Extern1	3200
OSPFv3 Extern2	3300
IS-IS L1 Ext	3400
IS-IS L2 Ext	3500

Usage Guidelines

Although these priorities can be changed, do not attempt any manipulation unless you are expertly familiar with the possible consequences. If you change the route priority, you must save the configuration and reboot the system.



Note

The priority for a blackhole route can not overlap with the priority of any other route origin.

Example

The following example sets the IPv6 route priority for static routing to 1200:

```
configure iproute ipv6 priority static 1200
```

History

This command was first available in ExtremeXOS 11.2.

The vr option was added in ExtremeXOS 12.1.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document..

configure iproute priority

```
configure iproute {ipv4} priority [auto-peering | blackhole | bootp
| ebgp | host-mobility | ibgp | icmp | isis | isis-level-1 | isis-
level-1-external | isis-level-2 | isis-level-2-external | mpls |
ospf-as-external | ospf-extern1 | ospf-extern2 | ospf-inter | ospf-
intra | rip | static | evpn | ospf-inter-vr] priority {vr vrname}
```

Description

Changes the priority for all routes from a particular route origin.

Syntax Description

auto-peering	Specifies the auto-peering route.
blackhole	Specifies the blackhole route.
bootp	Specifies BOOTP.
ebgp	Specifies E- <i>BGP</i> routes.
host-mobility	Host-Mobility route.
ibgp	Specifies I-BGP routes.
icmp	Specifies <i>ICMP</i> .

isis	Specifies IS-IS and applies only to blackhole routes installed for summary addresses.
isis-level-1	Specifies IS-IS Level 1 routing.
isis-level-1-external	Specifies IS-IS Level 1 External routing.
isis-level-2	Specifies IS-IS Level 2 routing.
isis-level-2-external	Specifies IS-IS Level 2 External routing.
mpls	Specifies <i>MPLS</i> routing.
ospf-as-external	Specifies <i>OSPF</i> as External 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 <i>RIP</i> .
static	Specifies static routes.
evpn	Specifies EVPN routes.
ospf-inter-vr	Specifies OSPF inter-VR routes.
<i>priority</i>	Specifies a priority number in the range of 11 to 65534.
<i>vrname</i>	Specifies a VR or VRF name.

Default

The following table lists the relative priorities assigned to routes depending upon the learned source of the route.

Table 12: Relative Route Priorities

Route Origin	Priority
Direct	10
MPLS	20
Blackhole	50
Static	1100
HostMobility	1150
ICMP	1200
EVPN	1698
Autopeering	1699
EBGP	1700
IBGP	1900
OSPFIntra	2200
OSPFInter	2300

Table 12: Relative Route Priorities (continued)

Route Origin	Priority
IS-IS	2350
IS-IS L1	2360
IS-IS L2	2370
RIP	2400
OSPFAsExt	3100
OSPF External 1	3200
OSPF External 2	3300
IS-IS L1 Ext	3400
IS-IS L2 Ext	3500
BOOTP	5000
OSPFInterVR	5010

Usage Guidelines

Although priorities can be changed, you should not attempt to do so unless you are certain of the possible consequences of the change. If you change the route priority, you must save the configuration and reboot the system.



Note

The priority for a blackhole route cannot overlap with the priority of any other route origin.

Example

The following example sets IP route priority for static routing to 1200:

```
# configure iproute priority static 1200
```

History

This command was first available in ExtremeXOS 10.1.

The route priority restrictions were added in ExtremeXOS 11.1.

The **ipv4** keyword was added in ExtremeXOS 11.2.

The **vr** option was added in ExtremeXOS 12.1.2.

The **evpn** option was added in ExtremeXOS 30.7.

The **ospf-inter-vr** option was added in ExtremeXOS 32.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure iproute reserved-entries

```
configure iproute reserved-entries [ num_routes_needed | maximum |
                                     default ] slot [all | slot_num]
```

Description

Reserves storage space for IPv4 and IPv6 routes in the Longest Prefix Match (LPM) hardware tables, allowing individual local and remote IPv4 unicast hosts to occupy the unused portions of the tables.

Syntax Description

<i>num_routes_needed</i>	Specifies a specific number of routes to reserve.
maximum	Reserves the maximum amount of space for IP route entries. No IPv4 hosts are stored in the LPM and External tables.
default	Reserves the default amount of space for IP route entries.
all	For SummitStack switches only, this option applies the reservation to all applicable slots.
<i>slot_num</i>	For SummitStack switches only, this option applies the reservation to the specified slot.

Usage Guidelines

Demand on the Layer 3 Hash table can be reduced by allowing IPv4 hosts to be stored in the LPM tables instead. This command allows you to reserve a portion of the LPM tables for routes, and this creates an unreserved portion that can be used to store IPv4 hosts. For more information, see the “Extended IPv4 Host Cache” section of the [Switch Engine 32.3 User Guide](#).

The default setting can support most networks, but if more than a few hundred local IP hosts and IP multicast entries are present, you can improve switch performance by calculating and configuring the reserved space for route entries to allow unreserved space for IPv4 hosts. Changing the number of reserved route entries does not require a reboot of the affected slots or switch.

You can view the current LPM hardware table usage by entering the `show iproute reserved-entries statistics` command. The LPM table statistics are in the columns under the In HW Route Table heading.

If the switch contains fewer routes than the capacity of the LPM tables, the number of route entries to reserve for a slot or switch should be the number of routes currently used in the hardware tables, plus an additional cushion for anticipated growth. Because each IPv6 route takes up the space of two IPv4 routes, the number of route entries to reserve is two times the value in the IPv6 routes column, plus the value in the IPv4 routes column, plus room for anticipated growth. For example, if you want to reserve space for 100 IPv4 routes and 20 IPv6 routes, the required number of route entries is 140 (100 + 2*20).

The maximum value for `num_routes_needed` for ExtremeSwitching 5420 switches is 12,256. For all other models, 16,352.

The maximum values shown above apply to ExtremeSwitching series switches operating independently or as part of a SummitStack. The maximum option can be used to specify the maximum values.

When maximum is specified, IPv4 hosts do not occupy LPM table space. Note that when maximum is specified, software forwarding can result, depending on the utilization and addresses in the Layer 3 Hash table, and is therefore not recommended.

When Algorithmic Longest-Prefix Match (ALPM) is configured using `configure forwarding internal-tables more routes`, the value for reserved-entries is treated as "maximum". Therefore, IPv4 hosts do not occupy LPM table space in order to maximize route capacity.

If the switch contains more routes than the capacity of the LPM tables, a trade-off can be made. You can choose to reserve 400 iproute entries, for example. The 400 IPv4 routes with the longest length network masks will be installed in the LPM table, and the remainder of the LPM table can be used for cache space for local and remote hosts. The remote host entries are only required for IPv4 addresses matching one of the 300 routes not installed in the LPM table. Since in this example, not all routes can be stored anyway, leaving appropriate room for individual remote hosts can result in more fast-path forwarding.

Depending on the actual routes present, IP route compression for IPv4 and/or IPv6 can be enabled to reduce the number of routes required in the LPM tables. For more information, see the description for the following command: `enable iproute compression {vr vrname}`

Example

The following command reserves up to 140 IPv4 routes or 70 IPv6 routes, or any combination in between, on all switches in a SummitStack:

```
# configure iproute reserved-entries 140 slot all
```

For details on the configuration changes, see the command descriptions for the following commands:

```
show iproute reserved-entries
```

```
show iproute reserved-entries statistics
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure iproute protection ping interval

```
configure iproute {ipv4 | ipv6} protection ping interval seconds miss
misses
```

Description

Configures the desired interval between pings and number of misses for ping protection of IPv4 and IPv6 static routes.

Syntax Description

ipv4	Designates IPv4 settings (default).
ipv6	Designates IPv6 settings.
protection	Configures route protection settings.
ping	Configures static route ping protection interval and number of misses.
interval	Number of seconds between pings to protected gateways. Ping response must be received within configured interval.
<i>seconds</i>	Number of seconds between pings to protected gateways. Range is 1–600. Default is 2.
miss	Number of pings with no response before associated routes are considered down.
<i>misses</i>	Number of pings with no response before associated routes are considered down. Range is 2–255. Default is 3.

Default

If not specified, IPv4 is the default, and:

- Interval = 2 seconds
- Misses = 3

Usage Guidelines

At the configurable interval, each unique gateway or device IP address configured for static route ping protection is sent an *ICMP* or ICMPv6 Echo Request if the ARP or Neighbor cache entry already has the IP->MAC binding. An ARP or Neighbor Solicitation is sent if the IP->MAC binding is unknown, and upon receiving a response, the ICMP Echo Request is sent.

The desired interval between pings and number of misses can be configured independently for IPv4 and IPv6.

Example

The following example sets for IPv4 a ping interval of 3 seconds and number of missed pings to 5:

```
# configure iproute ipv4 protection ping interval 3 miss 5
```

History

This command was first available in ExtremeXOS 22.1.

Platform Availability

This command is available on all platforms with any license level as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure iproute sharing hash-algorithm crc

```
configure iproute sharing {hash-method default} hash-algorithm crc
[lower | upper]
```

Description

This command is used to configure the "default" hash algorithm used to choose a gateway when hardware forwards an IPv4 or IPv6 unicast packet to a route with multiple equal-cost multipath gateways.

For information about configuring the custom hash method, see the command [configure iproute sharing hash-method custom](#).

The values within the IP unicast packet that are considered in the hash calculation depend on the setting of another command, `configure forwarding sharing [L3 | L3_L4]`. With the default, L3_L4, the hash calculation includes Source and Destination IP addresses, and the Source and Destination Layer 4 Port numbers. Or, if `configure forwarding sharing L3` is configured, the hash calculation only includes Source and Destination IP addresses. The distribution of packets among multiple gateways based on the IP Route Sharing lower or upper hash algorithm will depend on network traffic. The command will not result in traffic loss and takes effect immediately.

Syntax Description

iproute	IP routing module.
sharing	Configure settings for equal cost multipath routing";capability="route_sharing.
hash-method	Configures hardware forwarding hash method used to select among ECMP gateways for an IPv4 or IPv6 destination.
default	Default method for ECMP hardware hash calculation. For information about configuring the custom hash method, see the <code>configure iproute sharing hash-method custom {hash-algorithm [xor crc-16 crc-32 [lower upper]]}</code> command.
hash-algorithm	Configure hardware forwarding hash algorithm used to select among ECMP gateways for an IPv4 or IPv6 destination";capability="pib".
crc	Cyclic Redundancy Check (CRC).

lower	Lower bits of CRC32 hash calculation of source and destination packet criteria, used to select an ECMP gateway (Default).
upper	Upper bits of CRC32 hash calculation. May improve distribution when source and destination IP and ports do not vary much.

Default

Lower.

Usage Guidelines

Use this command to configure the hash algorithm used to choose a gateway when hardware forwards an IPv4 or IPv6 unicast packet to a route with multiple equal-cost multipath gateways. The values within the IP unicast packet that are considered in the hash calculation depend on the setting of another command, `configure forwarding sharing [L3 | L3_L4]`. With the default, `L3_L4`, the hash calculation includes Source and Destination IP addresses, and the Source and Destination Layer 4 Port numbers. Or, if `configure forwarding sharing L3` is configured, the hash calculation only includes Source and Destination IP addresses. The distribution of packets among multiple gateways based on the IP Route Sharing lower or upper hash algorithm will depend on network traffic. The command will not result in traffic loss and takes effect immediately.

Example

```
# configure iproute sharing hash-algorithm upper
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure iproute sharing max-gateways

```
configure iproute sharing max-gateways max_gateways
```

Description

Specifies the maximum number of gateways in each gateway set in the [ECMP](#) hardware table.

Syntax Description

<i>max_gateways</i>	Specifies the maximum number of ECMP gateways in a gateway. The only values allowed are 2, 4, 8, 16, 32 and 64. For ExtremeSwitching 5520 series switches, the allowed values are 2, 4, and 8.
---------------------	--

Default

For all platforms except the ExtremeSwitching 5520 series switches, the default is 16 gateways.

For ExtremeSwitching 5520 series switches, the default is 4.

Usage Guidelines

When IPv4 or IPv6 route sharing is enabled, the maximum number of gateways value represents the maximum number of next-hop gateways that can be used for communications with a destination subnet. Each gateway represents an alternative path to a subnet. The gateways can be defined with static routes, or they can be learned through the [OSPF](#), [OSPFv3](#), [BGP](#), or IS-IS protocols. The value for max-gateways applies to both IPv4 and IPv6 on all VRs.

When Pseudowire Label Switch Path Load Sharing is enabled, the maximum number of gateways value represents the maximum number of LSPs that a pseudowire can use for multi-path transport.

The max-gateways setting changes how the hardware is configured for multi-path; however, individual protocols have multi-path limitations that may be lower than the configured max-gateways setting. Additionally, the values supported for the max-gateways setting may vary, depending on the platform. See the [ExtremeXOS Release Notes](#) for the supported values of max-gateways for each protocol and platform.

The [ExtremeXOS Release Notes](#) also list the total number of route destinations and the total combinations of gateway sets that each platform can support with the different max-gateways option selections. For more information on selecting the maximum number of gateways and how this affects different platforms, see the “ECMP Hardware Table” in the [Switch Engine 32.3 User Guide](#).

You must save the configuration and reboot the switch for the new value to take effect. To see the current and configured value, use the commands `show ipconfig` or `show ipconfig ipv6`.

Example

The following example changes the maximum number of ECMP gateways per subnet or gateway set to 8:

```
configure iproute sharing max-gateways 8
```

History

This command was first available in ExtremeXOS 11.4.

The value 2 was first available in ExtremeXOS 12.0.2.

Support for shared gateway sets in the ECMP table was added in ExtremeXOS 12.4.

The values 16 and 32 were first available in ExtremeXOS 15.3.

This command first applied to IPv6 routes in ExtremeXOS 15.3.

The value 64 was added in ExtremeXOS 15.5.2

The default value for max. gateways was changed in ExtremeXOS 22.1 from 4 to 16. This applies only to new configurations. Existing configurations retain their settings.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches..

configure ip-security anomaly-protection icmp ipv4-max-size

```
configure ip-security anomaly-protection icmp ipv4-max-size size {slot
[ slot | all ]}
```

Description

Configures the maximum IPv4 *ICMP* allowed size.

Syntax Description

<i>size</i>	Specifies the size of the IPv4 ICMP in bytes.
<i>slot</i>	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default size is 512 bytes.

Usage Guidelines

This command configures the IPv4 ICMP allowed size. The absolute maximum is 1023 bytes.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ip-security anomaly-protection icmp ipv6-max-size

```
configure ip-security anomaly-protection icmp ipv6-max-size size {slot
[ slot | all ]}
```

Description

Configures the maximum ipv6 ICMP allowed size.

Syntax Description

<i>size</i>	Specifies the size of the IPv6 ICMP in bytes.
<i>slot</i>	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default size is 512 bytes.

Usage Guidelines

This command configures the IPv6 ICMP allowed size. The absolute maximum is 16K bytes.

You can use this command to configure the maximum IPv6 ICMP packet size for detecting IPv6 ICMP anomalies. If the next header in the IPv6 ICMP packet is not 0x3A:ICMP, this anomaly is not detected. For example, an IPv6 ICMP packet with packet header 0x2c: Fragment Header is not detected.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ip-security anomaly-protection notify cache

```
configure ip-security anomaly-protection notify cache size {slot [slot |  
all ]}
```

Description

Configures the size of local notification cache.

Syntax Description

<i>size</i>	Specifies the size of the local notification cache.
<i>slot</i>	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is 1000 events.

Usage Guidelines

This command configures the size of local notification cache. Cached events are stored in local memory. The range is between 1 and 1000 events per second. If the cache is full, newer events replace older events.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ip-security anomaly-protection notify rate limit

```
configure ip-security anomaly-protection notify rate limit value {slot
[slot | all ]}
```

Description

Configures the rate limiting for protocol anomaly notification.

Syntax Description

<i>value</i>	Specifies the period of the rate limit.
<i>slot</i>	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is 10 events per second.

Usage Guidelines

This is a paired command with `configure ip-security anomaly-protection notify rate window` that configures the rate limiting for protocol anomaly notification. When the anomaly notification is enabled, in order to avoid overloading CPU, the system generates only the number of limited notifications in a period of window seconds. The range is from 1 to 100 events.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ip-security anomaly-protection notify rate window

```
configure ip-security anomaly-protection notify rate window value {slot  
[slot | all ]}
```

Description

Configures the rate limiting for protocol anomaly notification.

Syntax Description

<i>value</i>	Specifies the period of the rate limit.
<i>slot</i>	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is 1 second.

Usage Guidelines

This is a paired command with `configure ip-security anomaly-protection notify rate limit` that configures the rate limiting for protocol anomaly notification. When the anomaly notification is enabled, in order to avoid overloading CPU, the system generates only the number of limited notifications in a period of window seconds. The range is between 1 and 300 seconds.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ip-security anomaly-protection notify trigger off

```
configure ip-security anomaly-protection notify trigger off value {slot  
[slot | all ]}
```

Description

Configures an anomaly rate-based notification feature.

Syntax Description

<i>value</i>	Specifies the number of events for the trigger.
<i>slot</i>	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is 1.

Usage Guidelines

This is a paired command with `configure ip-security anomaly-protection notify trigger on` that configures an anomaly rate-based notification feature. The anomaly notification is automatically triggered if the rate of anomaly events is greater than the configured ON value, and the notification is disabled if the rate falls below the value set in the `configure ip-security anomaly-protection notify trigger off` command.

The command takes effects after the anomaly notification is enabled.



Note

The value set in ON must be greater than or equal to the value set in OFF.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ip-security anomaly-protection notify trigger on

```
configure ip-security anomaly-protection notify trigger on value {slot
[slot | all ]}
```

Description

Configures an anomaly rate-based notification feature.

Syntax Description

<i>value</i>	Specifies the number of events for the trigger.
<i>slot</i>	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is 1.

Usage Guidelines

This is a paired command with `configure ip-security anomaly-protection notify trigger off` that configures an anomaly rate-based notification feature. The anomaly notification is automatically triggered if the rate of anomaly events is greater than the configured ON value, and the notification is disabled if the rate falls below the value set in the `configure ip-security anomaly-protection notify trigger off` command.

The command takes effects after the anomaly notification is enabled.



Note

The value set in ON must be greater than or equal to the value set in OFF.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ip-security anomaly-protection tcp

```
configure ip-security anomaly-protection tcp min-header-size size {slot
[ slot | all ]}
```

Description

Configures the minimum TCP header allowed.

Syntax Description

<i>size</i>	Specifies the size of the header in bytes.
<i>slot</i>	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default value is 20 bytes.

Usage Guidelines

This command configures the minimum TCP header allowed. It takes effect for both IPv4 and IPv6 TCP packets.

The range of the minimum TCP header may be between 8 and 255 bytes.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ip-security dhcp-bindings add

```
configure ip-security dhcp-binding add ip ip_address mac mac_address
  [dynamic vlan_id | {vlan} vlan_name] server-port server_port client-
port client_port lease-time seconds
```

Description

Creates a DHCP binding.

Syntax Description

<i>ip_address</i>	Specifies the IP address for the DHCP binding.
<i>mac_address</i>	Specifies the MAC address for the DHCP binding.
dynamic	Configuration options for dynamically created VLANs.
<i>vlan_id</i>	VLAN ID tag between 1 and 4,094.
<i>vlan_name</i>	Specifies the name of the <u>VLAN</u> for the DHCP binding.
<i>server_port</i>	Specifies the server port for the DHCP binding.
<i>client_port</i>	Specifies the client port for the DHCP binding.
<i>seconds</i>	Specifies the number of seconds for the lease.

Default

N/A.

Usage Guidelines

This commands allows you to add a DHCP binding in order to re-create the bindings after reboot and to allow IP Security features to work with clients having static IP addresses.



Note

Setting the lease-time to 0 causes the DHCP binding to be static; in other words, it is not aged-out if no DHCP renew occurs. This is for use with clients using static IP addresses.

History

This command was first available in ExtremeXOS 12.1.

Dynamic VLAN and VLAN ID options added in ExtremeXOS 30.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ip-security dhcp-bindings delete

```
configure ip-security dhcp-binding delete ip ip_address [dynamic vlan_id
| {vlan} vlan_name]
```

Description

Deletes a *DHCP* binding.

Syntax Description

<i>ip_address</i>	Specifies the IP address for the DHCP binding.
<i>vlan_name</i>	Specifies the name of the <i>VLAN</i> for the DHCP binding.
dynamic	Configuration options for dynamically created VLANs.
<i>vlan_id</i>	VLAN ID tag between 1 and 4,094.

Default

N/A.

Usage Guidelines

This commands allows you to delete a DHCP binding created with the command `configure ip-security dhcp-binding add ip ip_address mac mac_address {vlan} vlan_name server-port server_port client-port client_port lease-time seconds`.

History

This command was first available in ExtremeXOS 12.1.

Dynamic VLAN and VLAN ID options added in ExtremeXOS 30.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ip-security dhcp-bindings storage filename

```
configure ip-security dhcp-bindings storage filename name
```

Description

Creates a storage file for *DHCP* binding information.

Syntax Description

<i>name</i>	Specifies the name of the DHCP binding storage file.
-------------	--

Default

N/A.

Usage Guidelines

This commands allows you to configure the filename with which the DHCP bindings storage file is created on the external server when it is uploaded to the external server. The text file resides on an external server. You can configure the server with the command `configure ip-security dhcp-bindings storage location server [primary | secondary] ip_address | hostname {vrvr-name} tftp`.

The bindings file must have a .xsf extension. If the input filename doesn't already have a .xsf extension, one is added automatically.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ip-security dhcp-bindings storage location

```
configure ip-security dhcp-bindings storage location server [primary |  
secondary] ip_address | hostname]{vr vr-name} tftp
```

Description

Specifies the server location for the *DHCP* bindings storage file. The uploads can be made to any TFTP server regardless of the virtual router that it is present in.

Syntax Description

<i>ip_address</i>	Specifies the IP address location for the bindings storage file.
<i>hostname</i>	Specifies the hostname of the server.
<i>vr-name</i>	Specifies the virtual router name.
<i>none</i>	Using no option unconfigures the server.

Default

N/A.

Usage Guidelines

This commands allows you to specify where you want to store the DHCP storage file that you created with the command `configure ip-security dhcp-bindings storage filename name`.



Note

Using the command with no option unconfigures the server.

Example

The following command configures storage to the primary server 10.1.1.14:

```
configure ip-security dhcp-bindings storage location server primary 10.1.1.14 vr "VR-  
Default" tftp
```

The following example unconfigures the primary server:

```
configure ip-security dhcp-bindings storage location server primary
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ip-security dhcp-bindings storage

```
configure ip-security dhcp-bindings storage [write-interval minutes |  
write-threshold num_changed_entries]
```

Description

Configures DHCP bindings file storage upload variables.

Syntax Description

<i>minutes</i>	Specifies the number of minutes for the write interval.
<i>num_changed_entries</i>	Specifies the limit for the write threshold.

Default

The default write threshold is 50 entries; the default write interval is 30 minutes.

Usage Guidelines

This commands allows you to configure the upload variables for the DHCP bindings file that you created with the command `configure ip-security dhcp-bindings storage filename name` and specified the location of with the command `configure ip-security dhcp-bindings storage location server [primary | secondary] ip_address | hostname {vrvr-name} tftp`.

For redundancy, the DHCP bindings file is uploaded to both the primary and the secondary server. The failure of one upload (for example, due to a TFTP server timeout) does not affect the upload of any other.

When the maximum file size limit is reached, no additional DHCP bindings can be uploaded until one of the older bindings is removed.

The point at which DHCP bindings can be uploaded can be configured to work in one of the following ways:

- **Periodic upload:** Upload every N minutes, provided that DHCP bindings have changed since the last upload.
- **Upload based on number of yet-to-be uploaded entries:** Allows you to configure the maximum number of changed entries that are allowed to accumulate before being uploaded.

The write interval is configurable from 5 minutes to 1 day, with a default value of 30 minutes. The default value of the write threshold is 50 entries, with a minimum of 25 and maximum of 200.

Additions and deletions are considered changes, but updates are not, which means that DHCP renewals of existing leases are not counted.

By default, the write interval is in effect, but not the write-threshold. You may change whichever of these you wish by explicitly configuring the value.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ip-security dhcp-snooping information check

```
configure ip-security dhcp-snooping information check
```

Description

Enables the *DHCP* relay agent option (option 82) checking in the server-originated packets.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command enables the checking of the server-originated packets for the presence of option 82. In some instances, a DHCP server may not properly handle a DHCP request packet containing a relay agent option. Use this command to prevent DHCP reply packets with invalid or missing relay agent options from being forwarded to the client. With checking enabled, the following checks and actions are performed:

- When the option 82 is present in the packet, the MAC address specified in the remote-ID sub-option is the switch system MAC address. If the check fails, the packet is dropped.
- When option 82 is not present in the packet, the DHCP packet is forwarded with no modification.

To disable this check, use the following command:

```
unconfigure ip-security dhcp-snooping information check
```

Example

The following command enables DHCP relay agent option checking:

```
configure ip-security dhcp-snooping information check
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ip-security dhcp-snooping information circuit-id port-information port

```
configure ip-security dhcp-snooping information circuit-id port-
information port_info port port
```

Description

Configures the port information portion of the circuit ID.

Syntax Description

<i>port_info</i>	Specifies the circuit ID port information in the format of <i>VLAN Info - Port Info</i> ; maximum length is 32 bytes.
<i>port</i>	Specifies the port for which <i>DHCP</i> Snooping should be enabled.

Default

The default value is the ASCII representation of the ingress port's *SNMP* ifIndex.

Usage Guidelines

This command allows you to configure the port information portion of the circuit ID whose format is *vlan_info - port_info* for each port. The parameter *port_info* is a string of up to 32 bytes in length. When a specific value is not configured for port information, the *port_info* defaults to the ASCII representation of the ingress ports's SNMP ifIndex.

History

This command was first available in ExtremeXOS 12.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ip-security dhcp-snooping information circuit-id vlan-information

```
configure ip-security dhcp-snooping information circuit-id vlan-
information vlan_info [dynamic | {vlan} vlan_name | all]
```

Description

Configures the *VLAN* info portion of the circuit ID of a VLAN.

Syntax Description

<code>vlan_info</code>	Specifies the circuit ID VLAN information for each VLAN in the format of <code>VLAN Info-Port Info</code> ; maximum length is 32 bytes.
<code>vlan_name</code>	Specifies the VLAN for which <u>DHCP</u> should be enabled.
all	Specifies all VLANs.
dynamic	Configuration options for dynamically created VLANs.

Default

The default value is the ASCII representation of the ingress VLAN's ID.

Usage Guidelines

This command allows you to configure the VLAN information portion of the circuit ID of a VLAN. The VLAN info is a string of characters of up to 32 bytes in length, and is entered in the format of `VLAN InfoPort Info`. When a specific value is not configured for a VLAN, `vlan_info` defaults to the ASCII representation of the ingress VLAN's ID.

History

This command was first available in ExtremeXOS 12.1.

Dynamic VLAN option added in ExtremeXOS 30.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ip-security dhcp-snooping information option

```
configure ip-security dhcp-snooping information option
```

Description

Enables the DHCP relay agent option (option 82).

Syntax Description

This command has no arguments or variables.

Default

The default is unconfigured.

Usage Guidelines

This command enables the DHCP relay agent option (option 82), which is inserted into client-originated DHCP packets before they are forwarded to the server.

To disable the DHCP relay agent option (option 82), use the following command:

```
unconfigure ip-security dhcp-snooping information option
```

Example

The following command enable the DHCP relay agent option:

```
configure ip-security dhcp-snooping information information option
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ip-security dhcp-snooping information policy

```
configure ip-security dhcp-snooping information policy [drop | keep |  
replace]
```

Description

Configures the *DHCP* relay agent option (option 82) policy.

Syntax Description

drop	Specifies to drop the packet.
keep	Specifies to keep the existing option 82 information in place.
replace	Specifies to replace the existing data with the switch's own data.

Default

The default value is replace.

Usage Guidelines

Use this command to set a policy for the relay agent. Packets can be dropped, the option 82 information can be replaced (the default), or the packet can be forwarded with the information unchanged.

Example

The following command configures the DHCP relay agent option 82 policy to keep:

```
configure ip-security dhcp-snooping information information policy keep
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ip-security dhcp-snooping information remote-id

```
configure ip-security dhcp-snooping information remote-id [system-name |  
remote-id_info]
```

Description

Configures the *DHCP* relay agent remote ID.

Syntax Description

remote-id	Specifies configuring the remote ID.
system-name	Specifies assigning the switch's system name as the remote ID.
<i>remote-id_info</i>	Specifies assigning a user-defined string as the remote ID (up to 32 characters).

Default

If neither a system name nor the customized remote ID is configured, the default is the switch's MAC address.

Usage Guidelines

This command specifies setting the remote ID as either the switch's system name (for example, X465-48P) or a user-defined string. If neither selection has been made, or you unconfigure the remote ID (unconfigure `ip-security dhcp-snooping information remote-id`), the default remote ID is the switch's MAC address. However, this default (MAC address) name does not appear in the `show ip-security dhcp-snooping information remote-id` command.

Example

The following command configures the DHCP remote ID as the switch's system name::

```
# configure ip-security dhcp-snooping information remote-id system-name
```

The following command configures the DHCP remote ID as "mydhcp":

```
# configure ip-security dhcp-snooping information remote-id mydhcp
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ipv6 dad

```
configure ipv6 dad [off | on | {on} attempts max_solicitations] [{vr}  
  vr_name | vr all]
```

Description

Configures the operation of the duplicate address detection (DAD) feature on the specified VR.

Syntax Description

<i>max_solicitations</i>	Specifies the number of times the DAD feature tests for a duplicate address. The range is 1 to 10, and the default value is 1.
<i>vr_name</i>	Specifies a VR on which to enable this feature.

Default

DAD status: On on VR-Default.

Maximum solicitations: 1 for VR-Default.

If you do not specify a VR or VRF, the current VR context is used.

Usage Guidelines

When the DAD feature is enabled, the switch checks for duplicate IPv6 addresses on the specified VR when an IPv6 interface is initialized, or when a DAD check is initiated with a CLI command. After initialization, and when this feature is off, the switch does not start DAD checks.

Changes to the number of solicitations configuration take affect the next time the DAD check is run.

By default, this command applies to the current VR context, if no VR name is specified. If vr all is specified, the command applies to all user VRs and VR-Default.

The DAD feature does not run on loopback VLANs.

Example

The following command enables the DAD feature on all user VRs and VR-Default:

```
configure ipv6 dad on vr all
```

History

This command was first available in ExtremeXOS 12.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ipv6 hop-limit

```
configure ipv6 hop-limit hop_limit {dont-specify-in-ra} [{vr} vr_name |  
{vlan} vlan_name | vlan all}
```

Description

This command allows you to configure the IPv6 hop-limit. This hop-limit is used in all originated IPv6 packets, and (if router discovery is enabled) in outgoing Router Advertisement packets as well.

Syntax Description

<i>hop_limit</i>	Hop limit for all originated IPv6 packets, and the advertised hop-limit for Router Advertisements. Hop limit value is between 1 and 255. Default is 64.
dont-specify-in-ra	Sets the advertised hop-limit in Router Advertisements to zero.
vr	Virtual router.
vlan	<i>VLAN</i> .
all	All VLANs.

Default

64.

Usage Guidelines

Use this command to configure the IPv6 hop-limit. The hop-limit is used in all originated IPv6 packets, and (if router discovery is enabled) in outgoing Router Advertisement packets as well.

The 0 value is special and used only in outgoing Router Advertisements to convey to the receiving hosts that the router has not specified a hop-limit value to be used when originating IPv6 packets. This can be configured by specifying the optional **dont-specify-in-ra** keyword. The hop-limit can

be configured for a VLAN, all VLANs in a Virtual Router, or all VLANs in the system. By default, the hop-limit is configured for all vlans in the current Virtual Router context of the CLI.

Example

Example output not yet available and will be provided in a future release.

History

This command was first available in ExtremeXOS 15.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure iqagent http-proxy

```
configure iqagent http-proxy [ipaddress [fqdn | ip_address] port
    port_number | user user_name password [encrypted encrypted_password |
    password] | none]
```

Description

Configures the IQ Agent HTTP Proxy server IP and port, and defines the username and password if required.

Syntax Description

iqagent	Specifies configuring IQ Agent.
http-proxy	Specifies the HTTP proxy server that the IQ Agent uses.
ipaddress	Specifies the HTTP proxy server address.
<i>fqdn</i>	Specifies Fully-qualified domain name. Type is string. Range is 1-255.
<i>ip_address</i>	Specifies the dotted decimal IP Address.
port	Specifies the HTTP proxy port number.
<i>port_number</i>	Specifies the port number.
user	Specifies the HTTP proxy user name.
<i>user_name</i>	Specifies the user name. Type is string. Range is 1-63.
password	Specifies the HTTP proxy password.
encrypted	Specifies that the password is encrypted.
<i>encrypted_password</i>	Specifies the encrypted password (in Base64). Type is string. Range is 16-255.
<i>password</i>	Specifies the password (in ASCII). Type is string. Range is 1-63.
none	Specifies to clear all http-proxy configuration.

Default

N/A.

Usage Guidelines

To view IQ Agent information, use the command `show iqagent discovery`.

Example

The following example configures the HTTP proxy server at the address 10.51.3.163 on port 3128:

```
# configure iqagent http-proxy ipaddress 10.51.3.163 port 3128
```

The following example configures the HTTP proxy user "iqagent" with an encrypted password:

```
# configure iqagent http-proxy user iqagent password encrypted 35m5wuDryaqLrbQfZ5y4zw==
```

History

This command was first available in ExtremeXOS 31.3.

Platform Availability

This command is available on ExtremeSwitching 5420 and 5520 series switches.

configure iqagent server

```
configure iqagent server [vr [[vr-name | none] | vr_name vlan vlan-name]] | none] | ipaddress [fqdn | ip_address | none]]
```

Description

Configures the optional user-defined virtual router (VR) and address for the server for ExtremeCloud™ IQ Agent to connect to.

Syntax Description

iqagent	Specifies configuring IQ Agent.
server	Specifies configuring IQ server.
vr	Specifies selecting a VR for IQ server.
<i>vr-name</i>	Specifies the name of the VR.
none	Specifies no user-defined VR. Auto-discovery is used instead (default).
vlan	Specifies which VLAN to run on.
<i>vlan-name</i>	Specifies the name of the VLAN.
none	Specifies no user-defined VLAN.

ipaddress	Specifies configuring a server address
<i>fqdn</i>	Specifies a fully-qualified domain name for IQ server.
<i>ip_address</i>	Specifies an IP address for IQ server.
none	Specifies using auto-discovered information from IQ Agent discovery (default).

Default

By default, if no VR is specified, the VR is auto-discovered.

By default, for the server address, auto-discovery is used.

Usage Guidelines

To view IQ Agent information, use the command `show iqagent discovery`.

Example

The following example configures connecting to IQ server on VR "VR-Mgmt":

```
# configure iqagent server vr VR-Mgmt
```

The following example configures connecting to IQ server at the address 134.141.1.1:

```
# configure iqagent server ipaddress 134.141.1.1
```

The following example unconfigures any user-defined changes and returns to auto-discovery:

```
# configure iqagent server none
```

History

This command was first available in ExtremeXOS 30.7.

The **vlan** option was added in ExtremeXOS 32.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure irdp

```
configure irdp [multicast | broadcast | mininterval maxinterval lifetime preference]
```

Description

Configures the destination address of the router advertisement messages.

Syntax Description

multicast	Specifies multicast setting.
broadcast	Specifies broadcast setting.
<i>mininterval</i>	Specifies the minimum time between advertisements.
<i>maxinterval</i>	Specifies the maximum time between advertisements. Default is 600.
<i>lifetime</i>	Specifies the lifetime of the advertisement. Default is 1800.
<i>preference</i>	Specifies the router preference level. Default is 0.

Default

Broadcast (255.255.255.255). The default mininterval is 450.

Usage Guidelines

/CMP Router Discovery Protocol (IRDP) allows client machines to determine what default gateway address to use. The switch sends out IP packets at the specified intervals identifying itself as a default router. IRDP enabled client machines use this information to determine which gateway address to use for routing data packets to other networks.

Example

The following example sets the address of the router advertiser messages to multicast:

```
configure irdp multicast
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms that use the Edge, Advanced Edge, or Core license. For information on the licenses available for each platform, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure isis add vlan

```
configure isis add [vlan all | {vlan} vlan_name] area area_name {ipv4 | ipv6}
```

Description

This command associates the specified *VLAN* interface with the specified IS-IS router process.

Syntax Description

vlan all	Adds all IS-IS eligible VLANs to the router process.
<i>vlan_name</i>	Specifies a single IS-IS eligible VLAN to be added to the router process.
<i>area_name</i>	Identifies the router process to which the VLANs are added.
ipv4 ipv6	Specifies the VLAN IP address type, IPv4 or IPv6, to be added. If you do not specify an IP address type, the VLAN is added for the IPv4 address type. To support both IP address types on the same VLAN, enter the command twice, using a different IP address type each time.

Default

IPv4.

Usage Guidelines

An IS-IS-eligible interface is one that already has the appropriate IP address type (IPv4 or IPv6) address assigned to it. The VLAN must have an IPv4 address assigned to it if **ipv4** is specified or an IPv6 address assigned to it if **ipv6** is specified. In the event that a VLAN address is unconfigured, the interface is automatically removed from the IS-IS router.

VLANs are added to an IS-IS router process to form adjacencies with neighboring IS-IS routers. Hello PDUs are transmitted over these interfaces once the router process is enabled and has a system ID and area address. IP forwarding, IPv6 forwarding, or both must be enabled on the interface. If the router process operates at both L1 and L2, interfaces can be configured to form adjacencies in only a specific level.

Example

The following command adds VLAN SJvlan with an IPv4 address type to areax:

```
configure isis add SJvlan area areax
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Premier license as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure isis area add area-address

```
configure isis area area_name add area-address area_address
```

Description

This command adds an IS-IS area address to the specified routing process.

Syntax Description

<i>area_name</i>	Specifies the area name of the IS-IS process to which to add the area address.
<i>area_address</i>	Specifies an IS-IS area address to add to the IS-IS process. The area address can be from 1 to 13 bytes long and must be entered in the following format: 0101.0102.0103.0104.0105.0106.07.

Default

None.

Usage Guidelines

The IS-IS area address defines an L1 or L2 area within an AS. An IS-IS routing process must be assigned at least one area address before it can send or process PDUs. The area address must be configured appropriately. Level 1 routers only form adjacencies with other level 1 routers with at least one area address in common. Multiple area addresses may be configured, which may be desirable during a topological transition. The maximum number of area addresses that can be configured is 3.

Example

The following command assigns area address 0011.03 to areax:

```
configure isis area areax add area-address 0011.03
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Premier license as described in the [**Switch Engine 32.3 Feature License Requirements**](#) document.

configure isis area add summary-address

```
configure isis area area_name add summary-address [ipv4_address_mask |  
  ipv6_address_mask] {level [1 | 2]}
```

Description

This command adds an IPv4 or IPv6 summary address for the specified level on the specified router process.

Syntax Description

<i>area_name</i>	Specifies the router process to which the summary address is to be added.
<i>ipv4_address_mask</i>	Specifies an IPv4 summary address.
<i>ipv6_address_mask</i>	Specifies an IPv6 summary address.
level	Specifies the IS-IS level for the summary address. The level 1 option summarizes level 2 routes leaked to level 1. The level 2 option summarizes level 1 routes that are advertised into level 2.

Default

No summarization.

Usage Guidelines

Route summaries are useful for minimizing the number of LSPs required to describe reachability for an area. The summary address is advertised instead of the actual reachable addresses. This is particularly useful for L1/L2 routers in which the summary address is used in a single LSP instead of including a part or all of the addresses reachable in its level 1 area.

Note that a summary address is only advertised if at least one route matches the summary address. If there is no route present that matches the summary address exactly, a blackhole route is installed for the summary address. If an interlevel filter permits any route matched by the summary address, and that route is present, the summary address is advertised.

If multiple summary addresses are installed in which one or more supersede each other (10.0.0.0/8 and 10.0.0.0/16, for example), only the more specific summary addresses are advertised.

Example

The following command adds an IPv4 summary address to areax:

```
configure isis area areax add summary-address 10.0.0.0/8
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Premier license as described in the [**Switch Engine 32.3 Feature License Requirements**](#) document.

configure isis area area-password

```
configure isis area area_name area-password [none | [encrypted  
  simple encrypted_password | simple {password} ] {authenticate-snp  
  {tx-only}}]
```

Description

This command sets or clears the password for level 1 LSPs.

Syntax Description

<i>area_name</i>	Specifies the router process to which the password configuration applies.
none	Disables level 1 password authentication.
encrypted simple password	Enables password authentication and specifies that the supplied password is encrypted and must be decrypted prior to placement in a TLV.
authenticate-snp tx-only	Enables password authentication and level 1 SNP authentication. If the tx-only keyword is specified, the password is included in SNPs on transmission, but received SNPs are not authenticated.

Default

None.

Usage Guidelines

Only plain text passwords are supported. Passwords may be up to 254 alphanumeric characters in length. Although passwords are plaintext in the protocol, they are displayed and saved in an encrypted form.

When password authentication is enabled, received packets are authenticated against the configured password and are discarded if the password does not match. Authentication TLVs are included in transmitted level 1 LSPs with a configured password.

Example

The following command configures the password extreme for areax:

```
configure isis area areax area-password simple extreme
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Premier license as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure isis area delete area-address

```
configure isis area area_name delete area-address area_address
```

Description

This command deletes an area address from the specified routing process.

Syntax Description

area_name	Specifies the area name of the IS-IS process from which to delete the area address.
area_address	Specifies the area address name to delete from the IS-IS process.

Default

None.

Usage Guidelines

If this router process has only one area address configured, this command also causes the routing process to stop sending or processing IS-IS PDUs.

Example

The following command deletes the 0011.03 area address from areax:

```
configure isis area areax delete area-address 0011.03
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Premier license as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure isis area delete summary-address

```
configure isis area area_name delete summary-address [ipv4_address_mask
| ipv6_address_mask] {level [1 | 2]}
```

Description

This command removes the specified IPv4 or IPv6 summary address from the specified router process at the specified level.

Syntax Description

<i>area_name</i>	Specifies the router process from which the summary address is to be deleted.
<i>ipv4_address_mask</i>	Specifies an IPv4 summary address.
<i>ipv6_address_mask</i>	Specifies an IPv6 summary address.
level	Specifies the IS-IS level for the summary address.

Default

No summarization.

Usage Guidelines

Individual reachable addresses that were superseded by the summary address are now advertised in separate LSPs.

Example

The following command deletes an IPv4 summary address from areax:

```
configure isis area areax delete summary-address 10.0.0.0/8
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Premier license as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure isis area domain-password

```
configure isis area area_name domain-password [none | [encrypted  
simple encrypted_password | simple {password} ] {authenticate-snp  
{tx-only}]
```

Description

This command sets or clears the password for Level 2 LSPs.

Syntax Description

<i>area_name</i>	Specifies the router process for which the password is set or cleared.
none	Disables level 2 password authentication.
encrypted	Specifies that the supplied password is encrypted and must be decrypted prior to using it in a TLV.
<i>password</i>	Specifies a password. Passwords may be up to 254 alphanumeric characters in length.
authenticate-snp tx-only	If the optional authenticate-snp keyword is included, level 2 SNPs are also authenticated on receive and the password is included on transmission. If tx-only is specified, the password is included in SNPs on transmission, but received SNPs are not authenticated.

Default

None.

Usage Guidelines

Packets received are authenticated against the configured password and are discarded if the password does not match. Authentication TLVs are included in transmitted level 2 LSPs with the configured password. Only plain text passwords are supported. Although LSPs contain plain text passwords, passwords are displayed and saved in an encrypted form.

Example

The following command sets the domain password to Extreme:

```
configure isis area areax domain-password simple Extreme
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Premier license as described in the [**Switch Engine 32.3 Feature License Requirements**](#) document.

configure isis area interlevel-filter level 1-to-2

```
configure isis area area_name interlevel-filter level 1-to-2 [policy |  
    none] {ipv4 | ipv6}
```

Description

This command provides a method of restricting L1 routes from being redistributed into the L2 domain on an L1/L2 router.

Syntax Description

<i>area_name</i>	Specifies the router process for which this configuration change applies.
<i>policy</i>	Specifies a policy to control how L1 routes are redistributed.
none	Removes any previously configured interlevel filters.
ipv4 ipv6	Applies the interlevel filter to IPv4 or IPv6. If neither IPv4 nor IPv6 is specified, this command applies to IPv4.

Default

None.

Usage Guidelines

This command has no effect on level 1-only and level 2-only routers. Normally all L1 routes are redistributed into L2 on an L1/L2 router. Routes are permitted unless explicitly denied in the policy. This command does not necessarily disable level 1 to level 2 redistribution unless the configured policy effectively filters out all routes. For policies, the nlri match attribute is supported, and the permit and deny set attributes are supported.

Example

The following command removes any previously configured interlevel filters in areax for IPv4:

```
configure isis area areax interlevel-filter level 1-to-2 none
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Premier license as described in the [**Switch Engine 32.3 Feature License Requirements**](#) document.

configure isis area interlevel-filter level 2-to-1

```
configure isis area area_name interlevel-filter level 2-to-1 [policy |  
    block-all | allow-all] {ipv4 | ipv6}
```

Description

This command enables route leaking from level 2 to level 1 on an L1/L2 router.

Syntax Description

<i>area_name</i>	Specifies the router process for which this configuration change applies.
<i>policy</i>	Specifies a policy to control how L2 routes are leaked to L1.
block-all	Blocks all route leaking.
allow-all	Leaks all routes into level 1.
ipv4 ipv6	Applies the interlevel filter to IPv4 or IPv6. If neither IPv4 nor IPv6 is specified, this command applies to IPv4.

Default

block-all.

Usage Guidelines

When a policy is supplied with this command, all routes are leaked unless explicitly denied in the policy. This command has no effect on level 1-only and level 2-only routers. For policies, the nlri match attribute is supported, and the permit and deny set attributes are supported.

Example

The following command configures areax to leak all level 2 routes to level 1 for IPv4:

```
configure isis area areax interlevel-filter level 2-to-1 allow-all
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Premier license as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure isis area is-type level

```
configure isis area area_name is-type level [1 | 2 | both-1-and-2]
```

Description

This command configures the specified router process to operate as a level 1, level 2, or level 1/level 2 router.

Syntax Description

<i>area_name</i>	Specifies the router process you are configuring.
level	Specifies the IS-IS operation level for the router.

Default

both-1-and-2.

Usage Guidelines

Adjacencies are only formed with other routers of the same level. In addition, level 1 adjacencies are only formed with other level 1 routers with the same area address.

If there are no other L2 areas, the default is both-1-and-2. If an L2 or L1/L2 area is already present, the default is L1. This is because there can be only one L2 area in each system.

Example

The following command configures the areax router to operate at level 1:

```
configure isis area areax is-type level 1
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Premier license as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure isis area metric-style

```
configure isis area area_name metric-style [[narrow | wide]
{transition}] | transition] {level [1 | 2] }
```

Description

This command specifies the metric style for the specified router process and IS-IS level.

Syntax Description

<i>area_name</i>	Specifies the router process for which the metric style is to be configured.
narrow	Specifies the narrow metric style, which uses the 6-bit default metric. Only narrow metrics are encoded in originated TLVs; only narrow SPF calculations are performed.
narrow transition	Specifies the narrow metric style, which uses the 6-bit default metric. Only narrow metrics are encoded in originated TLVs; both narrow and wide SPF calculations are performed.
wide	Specifies the wide metric style, which uses the 24-bit metric specified in RFC 3784. Only wide metrics are encoded in originated TLVs; only wide SPF calculations are performed.
wide transition	Specifies the wide metric style, which uses the 24-bit metric specified in RFC 3784. Only wide metrics are encoded in originated TLVs; both narrow and wide SPF calculations are performed.
transition	Specifies both the narrow and wide metrics. Both narrow and wide metric types are encoded in TLVs; both narrow and wide SPF calculations are performed.
level	Specifies the IS-IS level to which the metric style applies.

Default

Narrow.

Usage Guidelines

Refer to RFC 3787, Section 5.1, for information on how to migrate a network from narrow metric-style to wide metric-style. Note that Section 5.2 is not supported. As a result, each interface's narrow and wide metric values must match while transitioning the metric style. Only when the entire network has transitioned to wide metric style should the interface metrics be configured differently than the configured narrow metric.

Example

The following command configures areax for the narrow metric style:

```
configure isis area areax metric-style narrow
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Premier license as described in the [**Switch Engine 32.3 Feature License Requirements**](#) document.

configure isis area overload-bit on-startup

```
configure isis area area_name overload-bit on-startup [ off | {suppress  
[external | interlevel | all]} seconds]
```

Description

This command enables or disables the overload bit feature while the specified IS-IS process is initializing.

Syntax Description

<i>area_name</i>	Specifies the area name of the IS-IS process for which this feature is to be enabled or disabled.
off	Disables the overload bit feature during initialization.
suppress	Specifies that one or all types of reachability information is to be suppressed or excluded from LSPs during initialization.
external	When included with the suppress option, this specifies that external reachability information is to be excluded from LSPs during initialization.
interlevel	When included with the suppress option, this specifies that interlevel reachability information is to be excluded from LSPs during initialization.
all	When included with the suppress option, this specifies that external and interlevel reachability information is to be excluded from LSPs during initialization.
<i>seconds</i>	Specifies the period (in seconds) during which this feature is enabled at initialization.

Default

Off.

Usage Guidelines

This command configures the overload bit to be set only while the configured router is initializing, and only for the period of time specified. This can be useful to minimize network churn while a new router joins and learns the topology. The suppress options are used during startup if the router process is level 1/level 2 or is running another protocol, such as *BGP* (in order to wait for the other protocol to

converge). Note that in the latter case, there is no signaling between protocols to indicate convergence. Again, this can reduce churn while the topologies are learned during router initialization.



Note

Although `enable isis area area_name overload-bit {suppress [external | interlevel | all]}` and `disable isis area area_name overload-bit` override the overload bit behavior configured by the `configure isis area area_name overload-bit on-startup [off | {suppress [external | interlevel | all]}seconds]` command, the enable and disable commands do not modify the configured parameters.

Example

The following command enables the areax overload bit feature for 15 seconds during initialization:

```
configure isis area areax overload-bit on-startup 15
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Premier license as described in the [**Switch Engine 32.3 Feature License Requirements**](#) document.

configure isis area system-id

```
configure isis area area_name system-id [automatic | system_id]
```

Description

This command configures the system ID for an IS-IS router process.

Syntax Description

<i>area_name</i>	Specifies the area name of the IS-IS process to which to add the system ID.
automatic	Sets the system ID to the system MAC address.
<i>system_id</i>	Specifies the 6-byte system ID using three sets of four hexadecimal digits, where each set is separated by a period. For example: 001B.1F62.1201.

Default

Automatic (system MAC address is used).

Usage Guidelines

The system ID must be a unique ID within the AS. Typically a system MAC address is used as the system ID. Sometimes a combination of one of the router's IP addresses and 2 prefix bytes are used. The assignment of the system ID may vary depending on how the AS is chosen to be administered.

Example

The following example configures an IS-IS system ID for areax:

```
configure isis area areax system-id 001B.1F62.1201
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Premier license as described in the [**Switch Engine 32.3 Feature License Requirements**](#) document.

configure isis area timer lsp-gen-interval

```
configure isis area area_name timer lsp-gen-interval seconds {level [1 | 2] }
```

Description

This command configures the minimum time required to wait before regenerating the same LSP.

Syntax Description

<i>area_name</i>	Specifies the router process for which you want to configure the LSP generation interval.
<i>seconds</i>	Specifies the generation level in seconds. The range is 1 to 120 seconds.
level	Specifies the level to which you want to apply the configuration. If neither level 1 nor level 2 is specified, the configuration applies to both levels.

Default

30 seconds.

Usage Guidelines

In link flapping situations in a mesh network, this can greatly reduce the amount of network traffic generated from LSP flooding.

Example

The following command sets the LSP generation interval to a value of 40 seconds:

```
configure isis area areax timer lsp-gen-interval 40
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Premier license as described in the [**Switch Engine 32.3 Feature License Requirements**](#) document.

configure isis area timer lsp-refresh-interval

```
configure isis area area_name timer lsp-refresh-interval seconds
```

Description

This command configures the refresh rate for locally originated LSPs.

Syntax Description

<i>area_name</i>	Specifies the router process for which you are setting the LSP refresh timer.
<i>seconds</i>	Specifies the LSP refresh interval. The range is 1 to 65535 seconds.

Default

900 seconds.

Usage Guidelines

This value should be configured to be less than the maximum LSP lifetime value, which is set with the `configure isis area area_name timer max-lsp-lifetimeseconds` command. Locally originated LSPs are purged and retransmitted at the specified interval regardless of link state.

Example

The following command sets the LSP refresh timer for areax to 1200 seconds:

```
configure isis area areax timer lsp-refresh-interval 1200
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Premier license as described in the [**Switch Engine 32.3 Feature License Requirements**](#) document.

configure isis area timer max-lsp-lifetime

```
configure isis area area_name timer max-lsp-lifetime seconds
```

Description

This command configures the LSP lifetime timer for locally originated LSPs.

Syntax Description

<i>area_name</i>	Specifies the router process for which you want to configure the LSP lifetime timer.
<i>seconds</i>	Specifies the LSP lifetime in seconds. The range is 1 to 65535 seconds.

Default

1200 seconds.

Usage Guidelines

This value should be configured to be greater than the LSP refresh interval, which is set with the `configure isis area area_name timer lsp-refresh-interval seconds` command. The remaining lifetime value is included in LSPs when they are flooded. Routers age out LSPs from other routers using the remaining lifetime provided in the LSP. If a refreshed version of the LSP is not received before it is aged out, an SPF recalculation occurs, possibly resulting in routing around the router from which the LSP originated.

Example

The following command configures the LSP lifetime timer for 1800 seconds:

```
configure isis area areax timer max-lsp-lifetime 1800
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Premier license as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure isis area timer restart

```
configure isis area area_name timer restart seconds {level [1 | 2]}
```

Description

This command configures the IS-IS T2 timer for the specified router process and level.

Syntax Description

area_name	Specifies the router process for which the T2 timer configuration applies.
seconds	Specifies the T2 timer value. The range is 5 to 65535 seconds.
level	Specifies the IS-IS level to which this timer configuration applies. If neither level 1 nor level 2 is specified, the configuration applies to both levels.

Default

60 seconds.

Usage Guidelines

The T2 timer is the restart timer for the LSP database for an IS-IS level. If the T2 timer for the respective level expires before the database has been resynchronized, SPF is run for that level.

Example

The following command configures the areax level 1 T2 timer for 90 seconds:

```
configure isis area areax timer restart 90 level 1
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Premier license as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure isis area timer spf-interval

```
configure isis area area_name timer spf-interval seconds {level [1 | 2] }
```

Description

This command specifies the minimum time to wait between SPF calculations.

Syntax Description

<i>area_name</i>	Specifies the router process for which you are configuring the SPF interval.
<i>seconds</i>	Specifies the minimum time between SPF calculations. The range is 1 to 120 seconds.
level	Specifies the IS-IS level to which the timer configuration applies. If neither level 1 nor level 2 is specified, the configuration applies to both levels.

Default

10 seconds.

Usage Guidelines

This helps prevent switch CPU overloading when a link flap causes several back-to-back SPF calculations.

Example

The following command configures the SPF interval timer for 30 seconds on areax:

```
configure isis area areax timer spf-interval 30
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Premier license as described in the [**Switch Engine 32.3 Feature License Requirements**](#) document.

configure isis area topology-mode

```
configure isis area area_name topology-mode [single | multi |  
transition] {level [1 | 2] }
```

Description

This command enables or disables use of multi-topology TLVs as specified in draft-ietf-isis-wg-multi-topology-11.

Syntax Description

<i>area_name</i>	Specifies the router process to be configured.
single	Specifies a single topology, where extended TLVs are used in SPF calculation and TLVs.
multi	Specifies a multi topology, where only the multi-topology TLVs are used in SPF calculation and TLVs.
transition	Specifies a transition topology, where both extended and multi-topology TLVs are used in SPF calculation and TLVs. The transition option is useful when migrating a routing domain.
level	For L1/L2 routers, this applies the configuration to IS-IS level 1 or level 2. If the level option is not specified, the configuration applies to both L1 and L2 areas. This option has no affect on L1-only and L2-only routers.

Default

Single.

Usage Guidelines

Multi-topology capability is desirable if both an IPv4 topology and an IPv6 topology exist with different routing paths.

Extreme supports MT IDs 0 and 2 (IPv4 unicast and IPv6 unicast) only.

Example

The following command configures the transition topology mode for areax:

```
configure isis area areax topology-mode transition
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Premier license as described in the [**Switch Engine 32.3 Feature License Requirements**](#) document.

configure isis circuit-type

```
configure isis [vlan all | {vlan} vlan_name] circuit-type level [1 | 2 | both-1-and-2]
```

Description

This command configures the circuit type level for one or all IS-IS VLANs.

Syntax Description

vlan all	Applies the selected circuit type to all IS-IS VLANs.
<i>vlan_name</i>	Specifies a single VLAN to which the circuit type configuration applies.
level [1 2 both-1-and-2]	Sets the circuit type level to level 1, level 2, or to both level 1 and level 2.

Default

Both-1-and-2.

Usage Guidelines

Hello PDUs are only sent on the specified level for the selected VLANs. This can be useful for level 1/level 2 routers that are neighbors.

Note that for per-level VLAN configurable parameters L1 and L1/L2, point-to-point interfaces use the level 1 parameters, and L2-only point-to-point interfaces use the L2 parameters.

Example

The following command configures all IS-IS VLANs to use circuit type level 1:

```
configure isis vlan all circuit-type level 1
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Premier license as described in the Switch Engine 32.3 Feature License Requirements document.

configure isis delete vlan

```
configure isis delete [vlan all | {vlan} vlan_name] {area area_name}  
                    {ipv4 | ipv6}
```

Description

This command removes a VLAN interface from the specified router process.

Syntax Description

vlan all	Deletes all IS-IS VLANs.
<i>vlan_name</i>	Specifies a single VLAN to delete.
<i>area_name</i>	Specifies the router process from which the VLAN is deleted. If you do not specify an IS-IS area, the software deletes the VLAN from the configured IS-IS area.
ipv4 ipv6	Specifies the IP address type for which the VLAN is deleted. If you do not specify an IP address type, the VLAN for the IPv4 address type is deleted. If the VLAN was added as IPv6, the ipv6 option must be used to remove the VLAN. If the VLAN was added as both IPv4 and IPv6, each VLAN IP address type must be deleted with a separate command.

Default

N/A.

Usage Guidelines

The associated adjacency is removed, causing the removal of the corresponding LSP if there is one, and causing an SPF recalculation if the router process is enabled. Hello PDUs are no longer sent on the specified interface. This command applies to IS-IS-enabled VLANs only.

Example

The following command deletes the IPv4 address type for all VLANs in areax:

```
configure isis delete vlan all area areax
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Premier license as described in the **Switch Engine 32.3 Feature License Requirements** document.

configure isis hello-multiplier

```
configure isis [vlan all | {vlan} vlan_name] hello-multiplier multiplier
               {level [1 | 2]}
```

Description

This command sets the hello multiplier for one or all IS-IS VLANs.

Syntax Description

vlan all	Applies the configuration to all IS-IS VLANs.
<i>vlan_name</i>	Specifies a single VLAN to which the configuration applies.
<i>multiplier</i>	Sets the hello multiplier. The range is 2 to 100.
level [1 2]	Limits the configuration to either level 1 or level 2. If neither level 1 nor level 2 is specified, the configuration applies to both levels.

Default

3.

Usage Guidelines

The hello multiplier is used in conjunction with the hello interval to compute the holding time. The holding time is included in hello PDUs and is calculated by multiplying the hello multiplier by the hello interval. If the hello interval is set to minimal, the holding time is set to 1 second and the hello interval is calculated by dividing 1 second by the hello multiplier. For example, a hello interval of minimal and a hello multiplier of 4 means that the hold interval is set to 250 ms (and the holding time to 1 second). The holding time tells the neighboring router how long to wait before declaring the sending router dead.

Example

The following command sets the SJvlan hello multiplier to 4:

```
configure isis SJvlan hello-multiplier 4
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Premier license as described in the **Switch Engine 32.3 Feature License Requirements** document.

configure isis import-policy

```
configure isis import-policy [policy-map | none]
```

Description

This command applies a policy map for routes imported to the FIB from all IS-IS router processes on this virtual router.

Syntax Description

<i>policy-map</i>	Specifies the policy to apply.
none	Removes any policies assigned to this virtual router.

Default

None.

Usage Guidelines

IS-IS policy files support the following policy match conditions:

- *nlri IPv4-address/mask-len IPv6-address/mask-len*
- *route-origin [isis-level-1 | isis-level-2 | isis-level-1-external | isis-level-2-external]*

IS-IS policy files support the following policy action statements:

- *cost*

Example

The following command applies the IS-IS policy *policy2* to the virtual router:

```
configure isis import-policy policy2
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Premier license as described in the [**Switch Engine 32.3 Feature License Requirements**](#) document.

configure isis link-type

```
configure isis [vlan all | {vlan} vlan_name] link-type [broadcast |  
point-to-point]
```

Description

This command specifies the link type for one or all IS-IS *VLANs*.

Syntax Description

vlan all	Applies the link type configuration to all IS-IS VLANs.
<i>vlan_name</i>	Specifies a single IS-IS VLAN to which the link type configuration is applied.
broadcast	Selects the broadcast link type for the specified VLANs.
point-to-point	Selects the point-to-point link type for the specified VLANs.

Default

Broadcast.

Usage Guidelines

On broadcast interfaces, a DIS is elected. There is no DIS election on point-to-point interfaces. If it is known that only two routers will be present on a physical network, it may be desirable to set their connecting interfaces to point-to-point mode. This reduces the overhead associated with DIS election and periodic CSNP transmissions and processing. In addition, if the adjacency is both level 1 and level 2, only one set of hello PDUs are sent on a point-to-point interface whereas hello PDUs are sent for both levels on broadcast interfaces. Interfaces in point-to-point mode must have an IP address assigned to them. Unnumbered interfaces are not supported.

For point-to-point interfaces, level 1 parameters apply to L1-only and L1/L2 interfaces. Level 2 parameters apply to L2-only point-to-point interfaces.

Example

The following command configures all IS-IS VLANs to use the broadcast link type:

```
configure isis vlan all link-type broadcast
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Premier license as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure isis mesh

```
configure isis [vlan all | {vlan} vlan_name] mesh [block-none | block-all | block-group group_id]
```

Description

This command configures LSP flooding behavior for the specified interface.

Syntax Description

vlan all	Applies the configuration to all IS-IS <i>VLANs</i> .
<i>vlan_name</i>	Specifies a single VLAN to which the configuration applies.
block-none	Disables LSP blocking.
block-all	Blocks all LSPs. No LSPs are flooded out of the selected interface.
block-group	Blocks LSPs that contain the specified group ID.
<i>group_id</i>	Specifies a group ID number. The range is 1 to 4294967295.

Default

Block-none.

Usage Guidelines

In a mesh environment, which is a set of fully interconnected point-to-point interfaces, LSP flooding can generate N2 PDUs because no router can tell which routers have and have not received the flooded LSP. By carefully selecting the links over which LSPs are flooded, traffic can be greatly reduced at the cost of some resiliency. Using mesh group IDs instead of a full block (the block-all option) allows a finer granularity of control.

Example

The following command configures blocking on SJvlan for group 5:

```
configure isis SJvlan mesh block-group 5
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Premier license as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure isis metric

```
configure isis [vlan all | {vlan} vlan_name] metric metric {level [1 | 2]}
```

Description

This command sets the narrow metric for one or all IS-IS [VLANs](#).

Syntax Description

vlan all	Applies the configuration to all IS-IS VLANs.
<i>vlan_name</i>	Specifies a single VLAN to which the configuration applies.
metric <i>metric</i>	Sets the metric value. The range is 1 to 63.
level [1 2]	Limits the configuration to either level 1 or level 2. If neither level 1 nor level 2 is specified, the configuration applies to both levels.

Default

10.

Usage Guidelines

If narrow metrics are enabled, this value is used in the associated LSPs for the selected VLANs.

Example

The following command sets the narrow metric for all IS-IS VLANs to 15:

```
configure isis vlan all metric 15
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Premier license as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure isis password vlan

```
configure isis [vlan all | {vlan} vlan_name] password [none | encrypted  
simple encrypted_password | simple { password }] [level [1|2]]
```

Description

This command sets or clears the authentication password for one or all IS-IS VLANs.

Syntax Description

vlan all	Applies the password configuration to all IS-IS VLANs.
<i>vlan_name</i>	Specifies a single VLAN to which the password configuration is applied.
none	Clears the password configuration and disables hello PDU authentication.
encrypted	Specifies that the supplied password is encrypted and must be decrypted prior to using it in a TLV.
<i>password</i>	Specifies the password. Passwords may be up to 254 alphanumeric characters in length.
level [1 2]	Limits the password configuration to level 1 or level 2. If neither level 1 or level 2 is specified, the configuration applies to both levels.

Default

None.

Usage Guidelines

If configured, the specified password is included in Hello PDUs for the specified level. In addition, received Hello PDUs on the specified interface are authenticated with the same password. Hello PDUs that are not authenticated are discarded.

Only plain text passwords are supported. Note that if the password is changed on an interface with an existing adjacency, the neighboring router needs to be configured as well. Depending on how timers are configured, the adjacency may time out while transitioning between passwords. Although passwords appear in plain text during configuration, they are displayed and saved in encrypted form.

Example

The following command assigns password Extreme to all level 1 VLANs configured for IS-IS:

```
configure isis vlan all password simple Extreme level 1
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Premier license as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure isis priority

```
configure isis [vlan all | {vlan} vlan_name] priority priority {level[1 | 2]}
```

Description

This command sets the priority used for DIS election on broadcast interfaces.

Syntax Description

vlan all	Applies the configuration to all IS-IS <i>VLANs</i> .
<i>vlan_name</i>	Specifies a single VLAN to which the configuration applies.
priority <i>priority</i>	Sets the priority value. The range is 0 to 127.
level [1 2]	Limits the configuration to either level 1 or level 2. If neither level 1 nor level 2 is specified, the configuration applies to both levels.

Default

64.

Usage Guidelines

A higher priority value is preferred over a lower priority value. The priority is encoded in level 1 or level 2 hello PDUs. This command is not valid for point-to-point interfaces. Note that a priority of 0 has no special meaning other than the fact that it is the lowest priority. A router with a priority of 0 can still become the DIS.

Example

The following command configures priority level 32 for SJvlan:

```
configure isis SJvlan priority 32
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Premier license as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure isis restart grace-period

configure isis restart grace-period *seconds*

Description

This command configures the T3 global restart timer for all IS-IS router processes on the current virtual router.

Syntax Description

<i>seconds</i>	Specifies the restart grace period in seconds. The range is 1 to 65535 seconds.
----------------	---

Default

65535.

Usage Guidelines

If the grace period expires before LSP resynchronization is complete, the virtual router sets the overload bit in LSPs that it originates.

Example

The following command sets the restart grace period to 5000 seconds:

```
configure isis restart grace-period 5000
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Premier license as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure isis restart

configure isis restart [**none** | **planned** | **unplanned** | **both**]

Description

This command configures IS-IS graceful restart behavior.

Syntax Description

none	Disables IS-IS graceful restart. When graceful restart is disabled, this router still operates as a helper to other restarting routers.
planned	Initiates IS-IS graceful restart only in response to the restart process isis.
unplanned	Initiates graceful restart only when the IS-IS process is restarted due to a process crash or an unplanned failover.
both	Initiates graceful restart for all events supported by the planned and unplanned options.

Default

None.

Usage Guidelines

The command options specify under which circumstances graceful restart is to be performed. This command has no effect during normal switch boot up. All IS-IS routing processes in the current virtual router are affected by this command.

All neighboring routers must support IS-IS restart in order for graceful restart to work. If graceful restart is not performed after a process restart or failover, the router's adjacencies are re-initialized causing SPF recalculation throughout the network and, if the overload bit is not configured to be set during startup, churn as adjacencies change state and LSPs are learned.



Note

The planned and unplanned command options do not affect the actual restart protocol operation of IS-IS; they only determine when the restart process occurs.

Example

The following command configures the switch to initiate a graceful restart for all events supported by the planned and unplanned options:

```
configure isis restart both
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Premier license as described in the [**Switch Engine 32.3 Feature License Requirements**](#) document.

configure isis timer csnp-interval

```
configure isis [vlan all | {vlan} vlan_name] timer csnp-interval seconds
               {level [1 | 2]}
```

Description

This command sets the minimum time between consecutive CSNP transmissions on the specified interface.

Syntax Description

vlan all	Applies the timer configuration to all IS-IS <i>VLANs</i> .
<i>vlan_name</i>	Specifies a single VLAN to which the timer configuration applies.
<i>seconds</i>	Sets the timer interval. The range is 1 to 65535 seconds.
level [1 2]	Limits the configuration to either level 1 or level 2. If neither level 1 nor level 2 is specified, the configuration applies to both levels.

Default

10 seconds.

Usage Guidelines

Periodic CSNPs are only sent on broadcast interfaces and only by the DIS.

Example

The following command sets the CSNP interval time for all IS-IS VLANs to 15 seconds:

```
configure isis vlan all timer csnp-interval 15
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Premier license as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure isis timer hello-interval

```
configure isis [vlan all | {vlan} vlan_name] timer hello-interval
               [seconds | minimal] {level [1 | 2]}
```

Description

This command sets the interval between two consecutive hello transmissions.

Syntax Description

vlan all	Applies the timer configuration to all IS-IS <i>VLANs</i> .
<i>vlan_name</i>	Specifies a single VLAN to which the timer configuration applies.
<i>seconds</i>	Sets the timer interval. The range is 1 to 65535 seconds.
minimal	Specifies that the hello interval is calculated by dividing 1 second by the hello multiplier.
level [1 2]	Limits the configuration to either level 1 or level 2. If neither level 1 nor level 2 is specified, the configuration applies to both levels.

Default

10 seconds.

Usage Guidelines

If this router is the elected DIS, hellos are sent three times more frequently than the configured interval.

When the timer configuration is set to minimal, the holding time included in the PDU is set to 1 second. Otherwise, the holding time is computed by multiplying the hello interval by the hello multiplier. The holding time tells the neighboring router how long to wait before declaring the sending router dead.

Example

The following command sets the hello interval timer for all VLANs to 15 seconds:

```
configure isis vlan all timer hello-interval 15
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Premier license as described in the [*Switch Engine 32.3 Feature License Requirements*](#) document.

configure isis timer lsp-interval

```
configure isis [vlan all | {vlan} vlan_name] timer lsp-interval  
milliseconds
```

Description

This command sets the minimum time between LSP transmissions.

Syntax Description

vlan all	Applies the timer configuration to all IS-IS <u>VLANs</u> .
<i>vlan_name</i>	Specifies a single VLAN to which the timer configuration applies.
<i>milliseconds</i>	Specifies the timer value. The range is 1 to 4294967295 milliseconds.

Default

33 milliseconds.

Usage Guidelines

This is used to throttle LSP flooding. Higher values reduce network traffic and can help keep underpowered routers from becoming overloaded during network events. Lower values speed up convergence.

Example

The following command sets the minimal LSP interval for IS-IS VLANs to 66 milliseconds:

```
configure isis vlan all timer lsp-interval 66
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Premier license as described in the Switch Engine 32.3 Feature License Requirements document.

configure isis timer restart-hello-interval

```
configure isis [vlan all | {vlan} vlan_name] timer restart-hello-  
interval seconds {level [1 | 2] }
```

Description

This command configures the T1 restart retransmit timer for one or all VLANs.

Syntax Description

vlan all	Specifies that the T1 restart timer configuration applies to all VLANs.
<i>vlan_name</i>	Specifies a VLAN to which the T1 restart timer configuration applies.
<i>seconds</i>	Specifies the T1 restart timer value. The range is 1 to 65535 seconds.
level [1 2]	Limits the configuration change to level 1 or level 2. If neither level 1 nor level 2 is specified, the configuration applies to both levels.

Default

3 seconds.

Usage Guidelines

If, after sending a restart request, the router process associated with this interface does not receive a restart acknowledgement and a CSNP within the period specified by this command, another restart request is sent.

Example

The following command sets the T1 restart timer to 6 seconds on all level 1 VLANs:

```
configure isis vlan all timer restart-hello-interval 6 level 1
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Premier license as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure isis timer retransmit-interval

```
configure isis [vlan all | {vlan} vlan_name] timer retransmit-interval  
seconds
```

Description

This command sets the time to wait for an acknowledgement of a transmitted LSP on a point-to-point interface.

Syntax Description

vlan all	Applies the timer value to all IS-IS <u>VLANs</u> .
<i>vlan_name</i>	Specifies a single VLAN to which the timer configuration applies.
<i>seconds</i>	Defines the timer value. The range is 0 to 65535 seconds.

Default

5 seconds.

Usage Guidelines

If an acknowledgement is not received when the timer expires, the LSP is resent and the timer is reset.

Example

The following command sets the retransmit interval for the SJvlan to 10 seconds:

```
configure isis SJvlan timer retransmit-interval 10
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Premier license as described in the [**Switch Engine 32.3 Feature License Requirements**](#) document.

configure isis wide-metric

```
configure isis [vlan all | {vlan} vlan_name] wide-metric metric {level[1 | 2]}
```

Description

This command sets the wide metric value for one or all IS-IS VLANs.

Syntax Description

vlan all	Applies the configuration to all IS-IS VLANs.
<i>vlan_name</i>	Specifies a single VLAN to which the configuration applies.
<i>metric</i>	Sets the metric. The range is 1 to 16777214.
level [1 2]	Limits the configuration change to either level 1 or level 2. If neither level 1 nor level 2 is specified, the configuration applies to both levels.

Default

10.

Usage Guidelines

If the wide metric style is enabled on the associated IS-IS router process, the wide metric value is used in Extended IP reachability TLVs, Extended IS Reachability TLVs, and IPv6 Reachability TLVs in LSPs.

Example

The following command sets the wide metric to 15 for all IS-IS VLANs:

```
configure isis vlan all wide-metric 15
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Premier license as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure jumbo-frame-size

```
configure jumbo-frame-size framesize
```

Description

Sets the maximum jumbo frame size for the switch.

Syntax Description

<i>framesize</i>	Specifies a maximum transmission unit (MTU) size for a jumbo frame. The range is 1523 to 9216; the default is 9216.
------------------	---

Default

Jumbo frames are disabled by default. The default size 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 framesize 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.

**Note**

Extreme Networks recommends that you set the MTU size so that fragmentation does not occur.

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.

Example

The following command configures the jumbo frame size to 5500:

```
configure jumbo-frame-size 5500
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure keychain accept-tolerance

```
configure keychain keychain_name accept-tolerance seconds
```

Description

This command configures a keychain accept tolerance in seconds.

Syntax Description

<i>keychain_name</i>	Specifies the name of the keychain.
accept-tolerance	Specifies the length of time an expired key can be accepted for received packets.
<i>seconds</i>	Specifies the tolerance in seconds. Range is 0-600 (default is 0).

Default

The default is 0, no tolerance.

Usage Guidelines

Use this command to configure a keychain accept tolerance.

Example

The following command configures the keychain accept tolerance:

```
configure keychain accept-tolerance 55
```

History

This command was first available in ExtremeXOS 31.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure keychain add key

```
configure keychain keychain_name add key key_id key-string [text_string
{active-lifetime {local} start start_time [end end_time | [duration
[seconds | maximum]]}] | encrypted encrypted_string]
```

Description

This command configures a key to add to a keychain.

Syntax Description

<i>keychain_name</i>	Specifies the name of the keychain.
add	Specifies adding a key to the keychain.
key	Authentication key entry.
<i>key_id</i>	Specifies the unique identifier within a keychain. Range is 1-65535.
key-string	Specifies the shared secret text string for the key.
<i>text_string</i>	Specifies the string. Range is 1-127.
active-lifetime	Specifies the time period the key will be active.
local	Specifies the time in the local time zone instead of UTC.
start	Specifies the start of the time period.
<i>start_time</i>	Specifies the date and time the key will become active. Format is YYYY-MM-DDThh:mm:ss.
end	Specifies the end of the time period.
<i>end_time</i>	Specifies the date and time the key will stop being active. Format is YYYY-MM-DDThh:mm:ss.
duration	Specifies the length of time the key will be active.

<i>seconds</i>	Specifies the duration in seconds. Range is 1-15552000.
maximum	Specifies that the key will be active for 180 days from the start time.
encrypted	Specifies the key string in encrypted form.
<i>encrypted_string</i>	Specifies the encrypted string.

Default

N/A

Usage Guidelines

A maximum of 8 keys can be added to a keychain.

The maximum length of a key string is 127 characters.

The maximum validity period of a key is 180 days.

Example

The following command configures the keychain to add an OSPFv3 key:

```
create keychain ospfv3-keys1
```

The following command adds a key string to the keychain:

```
configure keychain ospfv3-keys1 add key 3 key-string auth3
```

The following command adds additional options to the keychain:

```
configure keychain ospfv3-keys1 add key 1 key-string auth1 active-lifetime local start  
2021-06-01T00:00:00 end 2021-07-01T00:00:00
```

History

This command was first available in ExtremeXOS 31.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure keychain delete key

```
configure keychain keychain_name delete key key_id
```

Description

This command deletes a key to add to a keychain.

Syntax Description

<i>keychain_name</i>	Specifies the name of the keychain.
delete	Specifies deleting a key from the keychain.
key	Authentication key entry.
<i>key_id</i>	Specifies the unique identifier within a keychain. Range is 1-65535.

Default

N/A

Usage Guidelines

Use this command to delete a key from a configured keychain.

Example

The following command deletes a key:

```
configure keychain ospfv3-keys1 delete key 3
```

History

This command was first available in ExtremeXOS 31.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure keychain key active-lifetime

```
configure keychain keychain_name key key_id active-lifetime {local}  
    start start_time [end end_time | [duration [seconds | maximum]]]
```

Description

This command configures the time period a key will be active.

Syntax Description

<i>keychain_name</i>	Specifies the name of the keychain.
key	Authentication key entry.
<i>key_id</i>	Specifies the unique identifier within a keychain. Range is 1-65535.
active-lifetime	Specifies the time period the key will be active.
local	Specifies the time in the local time zone instead of UTC.

start	Specifies the start of the time period.
<i>start_time</i>	Specifies the date and time the key will become active. Format is YYYY-MM-DDThh:mm:ss.
end	Specifies the end of the time period.
<i>end_time</i>	Specifies the date and time the key will stop being active. Format is YYYY-MM-DDThh:mm:ss.
duration	Specifies the length of time the key will be active.
<i>seconds</i>	Specifies the duration in seconds. Range is 1-15552000.
maximum	Specifies that the key will be active for 180 days from the start time.

Default

N/A.

Usage Guidelines

Use this command to configure the time period a key will be active.

Example

The following command configures the keychain active lifetime:

```
configure keychain ospfv3-keys1 key 3 active-lifetime local start 2021-08-01T00:00:00 end
2021-09-01T00:00:00
```

History

This command was first available in ExtremeXOS 31.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure keychain key hash-algorithm

```
configure keychain keychain_name key key_id hash-algorithm algorithm
```

Description

This command configures a keychain key hash-algorithm.

Syntax Description

<i>keychain_name</i>	Specifies the name of the keychain.
key	Specifies the authentication key entry.

<i>key_id</i>	Specifies the unique identifier within a keychain. Range is 1-65535.
hash-algorithm	Specifies the hash algorithm to be used for the key.
<i>algorithm</i>	Specifies the supported algorithms (default is hmac-sha-256).

Default

The default algorithm is hmac-sha-256.

Usage Guidelines

Use this command to specify the hash algorithm to be used for the key.

Example

The following command configures the keychain hash algorithm:

```
configure keychain ospfv3-keys1 key 2 hash-algorithm hmac-sha-512
```

History

This command was first available in ExtremeXOS 31.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure l2pt profile add profile

```
configure l2pt profile profile_name add protocol filter filter_name
    {action [tunnel {cos cos} {dscp dscp_value {replace}} | encapsulate
    | none]}
```

Description

Adds an entry to an L2PT profile.

Syntax Description

profile <i>profile_name</i>	Specifies the profile that defines L2PT configuration for L2 protocols.
add protocol filter <i>filter_name</i>	Adds the specified Layer 2 protocol filter.
action	Specifies the action to perform on PDUs of the protocol (the default value is tunnel).
tunnel	Specifies to tunnel PDUs through the network.

cos <i>cos</i>	Specifies to override the class of service for tunneled PDUs, and specifies the class of service value to use for tunneling PDUs.
dscp	Specifies to set DSCP in the outer IP header when tunneled over VXLAN network.
<i>dscp_value</i>	Specifies the DSCP value to use in the outer IP header when the inner IP header is not present or when the replace is specified (default is 0). Range is 0-63.
replace	Specifies to replace the DSCP in the outer IP header with the configured value (default is do not replace, copy the inner DSCP to the outer header).
encapsulate	Specifies to encapsulate PDUs at egress, and decapsulate L2PT packets at ingress.
none	Specifies to not participate in tunneling for this protocol.

Default

Disabled.

The default *dscp_value* is 0.

The default **replace** is to not replace the DSCP.

Usage Guidelines

Use this command to add an entry to an L2PT profile.

Example

The following example adds an entry to *my_l2pt_prof* to tunnel protocols in "mylist" at cos 2:

```
configure l2pt profile my_l2pt_prof add protocol filter mylist action tunnel cos 2
```

The following example adds an entry to *my_l2pt_prof* to encapsulate/decapsulate protocols in "mylist":

```
configure l2pt profile my_l2pt_prof add protocol filter mylist action encapsulate
```

The following example adds an entry to *my_l2pt_prof* that is in use by 2 services:

```
configure l2pt profile my_l2pt_prof add protocol filter mylist
```

The following example configures a DSCP value to be set in the outer IP header when the inner DSCP is not present:

```
configure l2pt profile "vxlan_cdp" add protocol filter cdp action tunnel dscp 2
```

The following example copies the inner DSCP to the outer header when the inner DSCP is present:

```
configure l2pt profile "vxlan_cdp" add protocol filter cdp action tunnel dscp 3
```

The following example overrides the outer DSCP (even if the inner DSCP is present):

```
configure l2pt profile "vxlan_cdp" add protocol filter cdp action tunnel dscp 2 replace
```


History

This command was first available in ExtremeXOS 15.5.

Support for DSCP on VXLAN supported platforms was added in ExtremeXOS 31.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure l2pt profile delete profile

```
configure l2pt profile profile_name delete protocol filter filter_name
```

Description

Deletes an entry to an L2PT profile.

Syntax Description

profile <i>profile_name</i>	Specifies the profile that defines L2PT configuration for L2 protocols.
delete protocol filter <i>filter_name</i>	Deletes the specified Layer 2 protocol filter.

Default

Disabled.

Usage Guidelines

Use this command to delete an entry to an L2PT profile.

Example

The following example deletes the entry for "mylist" from my_l2pt_prof:

```
configure l2pt profile my_l2pt_prof delete protocol filter mylist
```

The following example deletes the entry entry for "mylist" from my_l2pt_prof that is in use by a service:

```
configure l2pt profile my_l2pt_prof delete protocol filter mylist
```

History

This command was first available in ExtremeXOS 15.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure l2vpn add peer

```
configure l2vpn [vppls vppls_name | vpws vpws_name] add peer ipaddress
                {{core {full-mesh | primary | secondary} | spoke}
```

Description

Configures a VPLS, H-VPLS, or VPWS peer for the node you are configuring.

Syntax Description

<i>vppls_name</i>	Specifies the VPLS for which you are configuring a peer.
<i>vpws_name</i>	Specifies the VPWS for which you are configuring a peer.
<i>ipaddress</i>	Specifies the IP address of the peer node.
core	Specifies that the peer is a core node. This option applies only to VPLS peers.
full-mesh	Specifies that the peer is a core full-mesh node. This is the default setting if neither the core or spoke options are specified. This option applies only to VPLS peers.
primary	Specifies that the peer is an H-VPLS core node and configures a primary H-VPLS connection to that core node. This option applies only to H-VPLS peers.
secondary	Specifies that the peer is an H-VPLS core node and configures a secondary H-VPLS connection to that core node. This option applies only to H-VPLS peers.
spoke	Specifies that the peer is a H-VPLS spoke node. This option applies only to H-VPLS peers.

Default

N/A.

Usage Guidelines

Each VPLS or H-VPLS node supports up to 64 peers, and each VPWS supports one peer. H-VPLS core nodes can peer with other core nodes and/or spoke nodes. H-VPLS spoke nodes can peer with core nodes but not with other spoke nodes.

VPLS core nodes must be configured in a full-mesh with other core nodes. Thus, all core nodes in the VPLS must have a configured PW to every other core node serving this VPLS. By default, the best LSP is chosen for the PW. The underlying LSP used by the PW can be configured by specifying the named LSP using the CLI command `configure l2vpn [vppls vppls_name | vpws vpws_name] peeripaddress [add | delete] mpls lsp_lsp_name.`

H-VPLS spoke nodes establish up to two point-to-point connections to peer with core nodes. If both primary and secondary peers are defined for a spoke node, the spoke node uses one of the peers for all communications. If both peers are available, the spoke node uses the connection to the primary peer.

If the primary peer connection fails, the spoke node uses the secondary peer. If the primary peer later recovers, the spoke node reverts back to using the primary peer.

VPWS nodes establish a point-to-point connection to one peer.

The `l2vpn` keyword was introduced in ExtremeXOS Release 12.4 and is required when configuring a VPWS peer. For backward compatibility, the `l2vpn` keyword is optional when configuring a VPLS peer. However, this keyword will be required in a future release, so we recommend that you use this keyword for new configurations and scripts.

Example

The following command adds a connection from the local core switch to the core switch at 1.1.1.202:

```
configure l2vpn vpls vpls1 add peer 1.1.1.202
```

The following command adds a connection from the local core switch to the spoke switch at 1.1.1.201:

```
configure l2vpn vpls vpls1 add peer 1.1.1.201 spoke
```

The following command adds a primary connection from the local spoke switch to the core switch at 1.1.1.203:

```
configure l2vpn vpls vpls1 add peer 1.1.1.203 core primary
```

The following command adds a VPWS connection from the local node to the peer switch at 1.1.1.204:

```
configure l2vpn vpws vpws1 add peer 1.1.1.204
```

History

This command was first available in ExtremeXOS 11.6.

Support for H-VPLS was first available in ExtremeXOS 12.1.

The **l2vpn** and **vpws** keywords were first available in ExtremeXOS 12.4.

Platform Availability

This command is available only on the platforms that support *MPLS* as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure l2vpn add service

```
configure l2vpn [vpls vpls_name | vpws vpws_name] add service [{vlan}  
vlan_name | {vman} vman_name]
```

Description

Adds a VLAN or VMAN service to a VPLS or VPWS.

Syntax Description

<i>vpls_name</i>	Identifies the VPLS interface within the switch (character string).
<i>vpws_name</i>	Identifies the VPWS interface within the switch (character string).
<i>vlan_name</i>	Logically binds the VLAN to the specified VPLS or VPWS.
<i>vman_name</i>	Adds the named VMAN to the VPLS or VPWS.

Default

N/A.

Usage Guidelines

Only one VLAN or VMAN can be configured per VPLS or VPWS.

When a VLAN service is added to a VPLS or VPWS, the VLAN ID is locally significant to the switch. Thus, each VLAN VPLS or VPWS interface within the Layer 2 VPN can have a different VLAN ID. This greatly simplifies VLAN ID coordination between metro network access points. Traffic may be switched locally between VLAN ports if more than one port is configured for the VLAN.

When a VMAN service has been configured for a VPLS or VPWS, the VMAN ID is locally significant to the switch. Thus, each VMAN VPLS or VPWS interface within the Layer 2 VPN can have a different VMAN ID, just like the VLAN service. The only difference is that the Layer 2 VPN overwrites the outer VMAN tag on Layer 2 VPN egress and leaves the inner VLAN tag unmodified. Because the inner VLAN tag is considered part of the customer packet data, the VMAN service can be used to emulate port-based services. This is accomplished by configuring the Layer 2 VPN to strip the 802.1Q tag from the tunneled packet. Since the switch inserts the VMAN tag when the packet is received and the 802.1Q tag is stripped before the packet is sent on the VPLS or VPWS PW, all packets received on ports that are members of the VMAN are transmitted unmodified across the Layer 2 VPN. The command `configure l2vpn [vpls vpls_name | vpws vpws_name] dot1q tag exclude` is used to configure the switch to strip the 802.1Q tag on the VPLS.

The `l2vpn` keyword was introduced in ExtremeXOS Release 12.4 and is required when adding a service to VPWS. For backward compatibility, the `l2vpn` keyword is optional when adding a service to VPLS. However, this keyword will be required in a future release, so we recommend that you use this keyword for new configurations and scripts.

Example

The example below adds a VLAN and a VMAN to the named VPLS:

```
configure l2vpn vpls myvpls add service vlan myvlan
configure l2vpn vpls myvpls add service vman myvman
```

The following example adds a VLAN and a VMAN to the named VPWS:

```
configure l2vpn vpws myvpws add service vlan vlan2
```

The following example adds a vman:

```
configure l2vpn vpws myvpws add service vman vman2
```

History

This command was first available in ExtremeXOS 11.6.

The l2vpn and vpws keywords were first available in ExtremeXOS 12.4.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure l2vpn delete peer

```
configure l2vpn [vpls vpls_name | vpws vpws_name] delete peer [ipaddress | all]
```

Description

Deletes the specified VPLS or VPWS peer.

Syntax Description

<i>vpls_name</i>	Identifies the VPLS within the switch (character string).
<i>vpws_name</i>	Identifies the VPWS within the switch (character string).
<i>ipaddress</i>	Specifies the IP address for the peer node that is the endpoint of the VC-LSP. This option applies only to VPLS peers.
all	Deletes all VPLS or VPWS peers. This option applies only to VPLS peers.

Default

N/A.

Usage Guidelines

When the VPLS or VPWS peer is deleted, VPN connectivity to the peer is terminated. The all keyword can be used to delete all peers associated with the specified Layer 2 VPN.

The l2vpn keyword was introduced in ExtremeXOS Release 12.4 and is required when deleting a VPWS peer. For backward compatibility, the l2vpn keyword is optional when deleting a VPLS peer. However,

this keyword will be required in a future release, so we recommend that you use this keyword for new configurations and scripts.

Example

The following example removes connectivity to 1.1.1.202 from VPLS1:

```
configure vpls vpls1 delete peer 1.1.1.202
```

History

This command was first available in ExtremeXOS 11.6.

The l2vpn and vpws keywords were first available in ExtremeXOS 12.4.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in the [Switch Engine 32.3 Feature License Requirements](#)

configure l2vpn delete service

```
configure l2vpn [vpls vpls_name | vpws vpws_name] delete service [{vlan}  
  vlan_name | {vman} vman_name]
```

Description

Deletes the specified [VLAN](#) or VMAN service from the specified Layer 2 VPN.

Syntax Description

<i>vpls_name</i>	Identifies the VPLS interface within the switch (character string).
<i>vpws_name</i>	Identifies the VPWS interface within the switch (character string).
<i>vlan_name</i>	Logically binds the VLAN to the specified VPLS.
<i>vman_name</i>	Adds the named VMAN to the VPLS.

Default

N/A.

Usage Guidelines

If there are no services configured for the VPLS or VPWS, all PWs within the Layer 2 VPN are terminated from the switch.

The l2vpn keyword was introduced in ExtremeXOS Release 12.4 and is required when deleting a service from a VPWS. For backward compatibility, the l2vpn keyword is optional when deleting a service from

a VPLS. However, this keyword will be required in a future release, so we recommend that you use this keyword for new configurations and scripts.

Example

The following example removes a service interface from a VPLS:

```
configure vpls vpls1 delete vman vman1
```

History

This command was first available in ExtremeXOS 11.6.

The l2vpn and vpws keywords were first available in ExtremeXOS 12.4.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure l2vpn health-check vccv

```
configure l2vpn [vpls [vpls_name | all] | vpws [vpws_name |  
all]] health-check vccv {interval interval_seconds} {fault-multiplier  
fault_multiplier_number}
```

Description

Configures the Virtual Circuit Connectivity Verification (VCCV) health check test and fault notification intervals for the specified VPLS or VPWS instance.

Syntax Description

<i>vpls_name</i>	Identifies the VPLS instance for which health check is to be configured.
<i>vpws_name</i>	Identifies the VPWS instance for which health check is to be configured.
all	Specifies that the configuration applies to all VPLS instances on the local node.
<i>interval_seconds</i>	Defines the interval between health check tests. The range is 1 to 10 seconds.
<i>fault_multiplier_number</i>	Specifies how long health check waits before a warning level message is logged. The wait period is the <i>interval_seconds</i> multiplied by the <i>fault_multiplier_number</i> . The <i>fault_multiplier_number</i> range is 2 to 6.

Default

Interval is 5 seconds.

Fault multiplier is 4.

Usage Guidelines

The VCCV health-check configuration parameters can be configured at anytime after the VPLS has been created.

The `show l2vpn {vpls {{vpls_name}} | vpws {{vpws_name}}} {peeripaddress} {detail} | summary` command displays the configured interval_seconds and fault-multiplier_number values for the VPLS or VPWS and the VCCV activity state.

The l2vpn keyword is introduced in ExtremeXOS Release 12.4 and is required when configuring health check for a VPWS. For backward compatibility, the l2vpn keyword is optional when configuring health check for a VPLS. However, this keyword will be required in a future release, so Extreme Networks recommends that you use this keyword for new configurations and scripts.

Example

The following command configures the health check feature on the VPLS instance myvpls:

```
configure vpls myvpls health-check vccv interval 10 fault-notification 40
```

History

This command was first available in ExtremeXOS 12.1.

The l2vpn and vpws keywords were first available in ExtremeXOS 12.4.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure l2vpn peer mpls lsp

```
configure l2vpn [vpls vpls_name | vpws vpws_name] peer ipaddress [add | delete] mpls lsp lsp_name
```

Description

Adds or deletes a named LSP as a specified PW for the specified Layer 2 VPN peer.

Syntax Description

<i>vpls_name</i>	Identifies the VPLS within the switch (character string).
<i>vpws_name</i>	Identifies the VPWS within the switch (character string).
<i>ipaddress</i>	Specifies the IP address for the peer node that is the endpoint of the PW-LSP. This option applies only to VPLS peers.

add	Permits addition of up to four RSVP-TE LSPs to the VPLS peer.
delete	Removes the LSP specified by the <code>lsp_name</code> parameter from the PW-LSP aggregation list.
<i>lsp_name</i>	Removes the specified lsp.

Default

N/A.

Usage Guidelines

If all the named LSPs are deleted from the configured Layer 2 VPN peer, VPLS or VPWS attempts to use the best-routed path LSP, if one exists. The delete portion of this command cannot be used to remove a named LSP that was selected by the switch as the best LSP. If no LSPs exist to the peer, Layer 2 VPN connectivity to the peer is lost. Currently, the VPLS or VPWS PW uses only one LSP.

The `l2vpn` keyword was introduced in ExtremeXOS Release 12.4 and is required when configuring a VPWS instance. For backward compatibility, the `l2vpn` keyword is optional when configuring a VPLS instance. However, this keyword will be required in a future release, so we recommend that you use this keyword for new configurations and scripts.

Example

The following examples add and remove a named LSP:

```
configure l2vpn vpls vpls1 peer 1.1.1.202 add mpls lsp "to-olympic4"
configure l2vpn vpls vpls1 peer 1.1.1.202 delete mpls lsp "to-olympic4"
```

The following example adds a named LSP for a VPWS peer:

```
configure l2vpn vpws vpws1 peer 1.1.1.203 add mpls lsp "to-olympic5"
```

History

This command was first available in ExtremeXOS 11.6.

The `l2vpn` and `vpws` keywords were first available in ExtremeXOS 12.4.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure l2vpn peer

```
configure l2vpn [vpls vpls_name ] peer ipaddress [limit-learning number
| unlimited-learning]
```

Description

Configures the maximum number of MAC SAs (Source Addresses) that can be learned for a given VPLS or VPWS peer.

Syntax Description

<i>vpls_name</i>	Identifies the VPLS within the switch (character string).
<i>ipaddress</i>	Specifies the IP address for the peer node that is the endpoint of the PW-LSP. This option applies only to VPLS peers.
limit-learning	Specifies a limit to the number of MAC SAs to be learned for the specified VPLS and peer.
<i>number</i>	The maximum number of MAC SAs that can be learned for the specified VPLS and peer.
unlimited-learning	Specifies no limit to the number of MAC SAs to be learned for the specified VPLS and peer.

Default

Unlimited.

Usage Guidelines

This parameter can only be modified when the specified VPLS is disabled. The unlimited-learning keyword can be used to specify that there is no limit. The default value is unlimited-learning.

The l2vpn keyword was introduced in ExtremeXOS Release 12.4 and is required when configuring a VPWS instance. For backward compatibility, the l2vpn keyword is optional when configuring a VPLS instance. However, this keyword will be required in a future release, so we recommend that you use this keyword for new configurations and scripts.

Example

The following example causes no more than 20 MAC addresses to be learned on VPLS1's PW to 1.1.1.202:

```
configure vpls vpls1 peer 1.1.1.202 limit-learning 20
```

History

This command was first available in ExtremeXOS 11.6.

The l2vpn and vpws keywords were first available in ExtremeXOS 12.4.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure l2vpn vpls add peer ipaddress

```
configure {l2vpn} vpls vpls_name add peer ipaddress { static-pw
transmit-label outgoing_pw_label receive-label incoming_pw_label }
```

Description

Configures L2VPN VPLS service over MPLS Static PW.

Syntax Description

<i>vpls_name</i>	Specifies the VPLS for which you are configuring a peer.
<i>ipaddress</i>	Specifies the IP address of the peer node.
static-pw transmit-label <i>label</i>	Specifies the static pseudowire transmit label.
<i>outgoing_pw_label</i>	Specifies the name of the egress label.
receive-label	Specifies the static PW receive label.
<i>incoming_pw_label</i>	Specifies the name of the ingress label.

Default

N/A.

Usage Guidelines

Use this command to statically configure a new MPLS Ethernet PW for the specified VPLS. You must specify the outgoing (MPLS ingress) and incoming (MPLS egress) PW labels. Similarly, you must configure the peer with a static PW that has the reverse PW label mappings.

Locally, the *incoming_pw_label* must be unique and is allocated out of the static label space. The *outgoing_pw_label* must match the peer's configured incoming PW label.

Just like a signaled PW, a static PW can optionally be configured to use any type of tunnel LSP: LDP, RSVP-TE, or Static. In the case of RSVP-TE and LDP, those protocols must be configured and enabled and an LSP must be established before traffic can be transmitted over the static PW.

For Static LSPs, only the MPLS ingress LSP (or outgoing LSP) is specified. Unlike signaled PWs, there is no end-to-end PW communication that is used to verify that the PW endpoint is operational, and in the case of static LSPs, that the data path to the PW endpoint is viable. In the event of a network fault, if a secondary RSVP-TE LSP is configured or the routing topology changes such that there is an alternate LDP LSP, the static PW will automatically switch LSPs in order to maintain connectivity with the PW endpoint. Static LSPs can be protected proactively by configuring BFD to verify the static LSPs IP next hop connectivity. Optionally, the underlying LSP for the PW can be explicitly specified using a named LSP. When a named LSP is explicitly specified, only the specified named LSP is used to carry the PW. In the event that a specified named LSP is withdrawn, the VPLS/VPWS remains operationally down until the named LSP is restored.

Since VC Status signaling is not supported, the VC Status “standby” bit cannot be used to allow support for PW redundancy and H-VPLS. Consequently, only “core full-mesh” PWs are allowed to have statically configured labels.

Example

The following command configures a new MPLS ethernet pseudowire for vpls1 :

```
configure vpls vpls1 add peer 1.1.1.202
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on all platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure vpls add service

```
configure vpls vpls_name add service [{vlan} vlan_name | {vman}
vman_name]
```



Note

This command has been replaced with the following command: `configure l2vpn [vpls vpls_name | vpws vpws_name] add service [{vlan} vlan_name | {vman} vman_name]` . This command is still supported for backward compatibility, but it will be removed from a future release, so we recommend that you start using the new command.

Description

Configures service for VPLS.

Syntax Description

<i>vpls_name</i>	Identifies the VPLS interface within the switch (character string)
<i>vlan_name</i>	Logically binds the <u>VLAN</u> to the specified VPLS.
<i>vman_name</i>	Adds the named VMAN to the VPLS.

Default

N/A.

Usage Guidelines

This command configures the VPLS service for the specified *vpls_name*. The VPLS service may be a customer VLAN or a customer VMAN. Specifying the *vlan_name* logically binds the VLAN to the specified VPLS. Only one VLAN or VMAN may be configured per VPLS.

When a VLAN service has been configured for a VPLS, the VLAN is added to the VPLS specified by the *vpls_name*. The VLAN ID is locally significant to the switch. Thus, each VLAN VPLS interface within the VPLS network may have a different VLAN ID service bound to the VPLS. This greatly simplifies VLAN ID coordination between metro network access points. Traffic may be switched locally between VLAN ports if more than one port is configured for the VLAN.

When a VMAN service has been configured for a VPLS, the VMAN is added to the VPLS specified by *vpls_name*. The VMAN ID is locally significant to the switch. Thus, each VMAN VPLS interface within the VPLS network may have a different VMAN ID, just like the VLAN service. The only difference is that the VPLS network overwrites the outer VMAN tag on VPLS egress and leaves the inner VLAN tag unmodified. Because the inner VLAN tag is considered part of the customer packet data, the VMAN service can be used to emulate port-based services. This is accomplished by configuring the VPLS to strip the 802.1Q tag from the tunneled packet. Since the switch inserts the VMAN tag when the packet is received and the 802.1Q tag is stripped before the packet is sent on the VPLS PW, all packets received on ports that are members of the VMAN are transmitted unmodified across the VPLS. The command `configure vpls vpls_name dot1q tag exclude` is used to configure the switch to strip the 802.1Q tag on the VPLS.

Example

The example below adds a VLAN and a VMAN to the named VPLS:

```
configure vpls myvpls add service vlan myvlan
configure vpls myvpls add service vman myvman
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure l2vpn vpls peer static-pw

```
configure l2vpn vpls vpls_name peer ipaddress static-pw {transmit-label
    outgoing_pw_label receive-label incoming_pw_label }
```

Description

Changes the labels of a statically configured Ethernet PW for a VPLS.

Syntax Description

<i>vpls_name</i>	Specifies the VPLS for which you are configuring a peer.
peer	Specifies the peer IP address.
<i>ipaddress</i>	Specifies the IP address of the peer node.
static-pw	Specifies the static pseudowire.
transmit label	Specifies the pseudowire transmit label.
<i>outgoing_pw_label</i>	Specifies the name of the egress label.
receive-label	Specifies the static pseudowire receive label.
<i>incoming_pw_label</i>	Specifies the name of the ingress label.

Default

N/A.

Usage Guidelines

Use this command to change the labels of a statically configured Ethernet PW for a VPLS that already exists. Either or both the outgoing (*MPLS* ingress) and incoming (MPLS egress) PW labels can be specified. The peer must be similarly configured with a static PW that has the reverse PW label mappings. Locally, the *incoming_pw_label* must be unique and is allocated out of the static label space. The *outgoing_pw_label* must match the peer's configured incoming PW label. The L2VPN can remain operational during the change; however, the PW goes down and comes back up.

Example

The following command changes the VPLS label to "VPLS1":

```
# configure l2vpn vpls vpls1 peer static-pw 1.1.1.202
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on all platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure l2vpn vpls redundancy

```
configure {l2vpn}vpls vpls_name redundancy [esrp esrpDomain | eaps |
      stp]
```

Description

Configures a VPLS instance to provide protected access using the [EAPS](#) redundancy type, the specified [ESRP](#) domain, or [STP](#).

Syntax Description

<i>vpls_name</i>	Specifies the VPLS for which you are configuring protection.
<i>esrpDomain</i>	Configures a VPLS instance to provide protected access using the specified ESRP domain.
<i>eaps</i>	Configures a VPLS instance to use the EAPS redundancy type.
<i>stp</i>	Configures a VPLS instance to request an FDB relearning process on an adjacent node when STP responds to a topology change for a VLAN .

Default

Redundancy disabled.

Usage Guidelines

Only one redundancy mode can be configured at a time on a VPLS, and the VPLS must be disabled when the redundancy mode is configured. If you attempt to configure a second mode, an error appears. The current redundancy mode must be unconfigured before you configure a different redundancy mode.

The ESRP domain specified must be a valid ESRP domain of type vpls-redundancy. If not, the command is rejected with an appropriate error message. When a VPLS instance is associated with an ESRP domain, the user cannot delete the ESRP domain unless the VPLS redundancy has been unconfigured. For VPLS access protection to become fully functional, VPLS redundancy must also be configured on a second VPLS peer using the same VPLS name and ESRP domain.

Specify the redundancy type as EAPS when using redundant EAPS access rings. This configuration requires EAPS shared links to be configured between redundant VPLS nodes. This configures VPLS to use a PW between VPLS attachment nodes instead of using a customer VLAN. This configuration is only required when there is an EAPS ring on the VPLS service VLAN.



Note

The EAPS master should not be on a VPLS node.

The STP option enables VPLS interfaces to respond appropriately to STP topology changes in a VLAN. For example, if STP detects a link failure, it will flush the appropriate FDB entries to initiate relearning on the STP protected interfaces. When this option is selected and STP initiates relearning, the VPLS interfaces on the same VLAN also initiate relearning so that a new VLAN path to the VPLS core can be learned. For more information, including limitations and restrictions, see the “VPLS STP Redundancy Overview” Section in the [Switch Engine 32.3 User Guide](#).

The l2vpn keyword was introduced in ExtremeXOS Release 12.4. For backward compatibility, the l2vpn keyword is optional when configuring a VPLS instance. However, this keyword will be required in a future release, so we recommend that you use this keyword for new configurations and scripts.

Example

The following command adds redundancy to the vpls1 VPLS using the esrp1 domain:

```
configure l2vpn vpls vpls1 redundancy esrp esrp1
```

The following command specifies the EAPS redundancy type for the vpls2 VPLS:

```
configure l2vpn vpls vpls2 redundancy eaps
```

The following command specifies the STP redundancy type for the vpls3 VPLS:

```
configure l2vpn vpls vpls3 redundancy STP
```

History

This command was first available in ExtremeXOS 12.1.

The l2vpn keyword and the STP option were added in ExtremeXOS 12.4.

Platform Availability

This command is available only on the platforms that support *MPLS* as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure l2vpn vpws add peer ipaddress

```
configure l2vpn vpws vpws_name add peer ipaddress ipaddress {static-pw  
transmit-label outgoing_pw_label receive-label incoming_pw_label }
```

Description

Configures L2VPN VPWS service over *MPLS* Static PW.

Syntax Description

<i>vpws_name</i>	Specifies the VPWS for which you are configuring a peer.
ipaddress	Specifies the peer IP address.
<i>ipaddress</i>	Specifies the IP address of the peer node.
static-pw transmit-label	Specifies the static pseudowire transmit label.
<i>outgoing_pw_label</i>	Specifies the name of the egress label.
receive-label	Specifies the static PW receive label.
<i>incoming_pw_label</i>	Specifies the name of the ingress label.

Default

N/A.

Usage Guidelines

Use this command to statically configure a new MPLS Ethernet PW for the specified VPWS. You must specify the outgoing (MPLS ingress) and incoming (MPLS egress) PW labels. Similarly, you must configure the peer with a static PW that has the reverse PW label mappings.

Locally, the *incoming_pw_label* must be unique and is allocated out of the static label space. The *outgoing_pw_label* must match the peer's configured incoming PW label.

Just like a signaled PW, a static PW can optionally be configured to use any type of tunnel LSP: LDP, RSVP-TE, or Static. In the case of RSVP-TE and LDP, those protocols must be configured and enabled and an LSP must be established before traffic can be transmitted over the static PW.

For Static LSPs, only the MPLS ingress LSP (or outgoing LSP) is specified. Unlike signaled PWs, there is no end-to-end PW communication that is used to verify that the PW endpoint is operational, and in the case of static LSPs, that the data path to the PW endpoint is viable. In the event of a network fault, if a secondary RSVP-TE LSP is configured or the routing topology changes such that there is an alternate LDP LSP, the static PW will automatically switch LSPs in order to maintain connectivity with the PW endpoint. Static LSPs can be protected proactively by configuring BFD to verify the static LSPs IP next hop connectivity. Optionally, the underlying LSP for the PW can be explicitly specified using a named LSP. When a named LSP is explicitly specified, only the specified named LSP is used to carry the PW. In the event that a specified named LSP is withdrawn, the VPLS/VPWS remains operationally down until the named LSP is restored.

Since VC Status signaling is not supported, the VC Status “standby” bit cannot be used to allow support for PW redundancy and H-VPLS. Consequently, only “core full-mesh” PWs are allowed to have statically configured labels.

Example

The following command configures VPWS service for VPWS1 on peer 1.1.1.202:

```
configure vpws vpws1 add peer 1.1.1.202
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on all platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure l2vpn vpws peer static-pw

```
configure l2vpn vpws vpws_name peer ipaddress static-pw {transmit-label
    outgoing_pw_label receive-label incoming_pw_label }
```

Description

Changes the labels of a statically configured Ethernet pseudowire for a VPWS.

Syntax Description

<i>vpws_name</i>	Specifies the VPWS for which you are configuring a peer.
peer	Specifies the peer IP address.
<i>ipaddress</i>	Specifies the IP address of the peer node.
static-pw	Specifies the static pseudowire.
transmit label	Specifies the pseudowire transmit label.
<i>outgoing_pw_label</i>	Specifies the name of the egress label.
receive-label	Specifies the static pseudowire receive label.
<i>incoming_pw_label</i>	Specifies the name of the ingress label.

Default

N/A.

Usage Guidelines

Use this command to change the labels of a statically configured Ethernet pseudowire for a VPWS that already exists. Either or both the outgoing (MPLS ingress) and incoming (MPLS egress) PW labels can be specified. The peer must be similarly configured with a static PW that has the reverse PW label mappings. Locally, the *incoming_pw_label* must be unique and is allocated out of the static label space. The *outgoing_pw_label* must match the peer's configured incoming PW label. The L2VPN can remain operational during the change; however, the PW goes down and comes back up.

Example

The following command changes the VPWS label to "vpws1":

```
# configure l2vpn vpws vpws1 peer static-pw 1.1.1.202
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on all platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure l2vpn

```
configure l2vpn [vpls vpls_name | vpws vpws_name] {dot1q [ethertype
    hex_number | tag [include | exclude]]} {mtu number}
```

Description

Configures VPLS or VPWS parameters.

Syntax Description

<i>vpls_name</i>	Identifies the VPLS within the switch (character string).
<i>vpws_name</i>	Identifies the VPWS within the switch (character string).
dot1q	Specifies the action the switch performs with respect to the 802.1Q ethertype or tag.
ethertype	Overwrites the ethertype value for the customer traffic sent across the PW
<i>hex_number</i>	Identifies the ethertype, uses the format of 0xN.
tag	Specifies the action the switch performs with respect to the 802.1Q tag.
include	Includes the 802.1Q tag when sending packets over the VPLS L2 VPN.
exclude	Strips the 802.1Q tag before sending packets over the VPLS L2 VPN.
mtu	Specifies the MTU value of the VPLS transport payload packet.
<i>number</i>	The size (in bytes) of the MTU value. The configurable MTU range is 1492 through 9216. The default VPLS MTU value is 1500.

Default

dot1q tag - excluded.

ethertype - the configured switch ethertype is used.

number (MTU) - 1500.

Usage Guidelines

This command configures the VPLS and VPWS parameters. PWs are point-to-point links used to carry VPN traffic between two devices within the VPLS. Each device must be configured such that packets transmitted between the endpoints are interpreted and forwarded to the local service correctly. The optional ethertype keyword may be used to overwrite the Ethertype value for the customer traffic sent across the PW. By default, the configured switch ethertype is used. If configured, the ethertype in the outer 802.1q field of the customer packet is overwritten using the configured ethertype value. The ethertype value is ignored on receipt.

Optionally, the switch can be configured to strip the 802.1q tag before sending packets over the VPLS or VPWS Layer 2 VPN. This capability may be required to provide interoperability with other vendor products or to emulate port mode services. The default configuration is to include the 802.1q tag.

The `mtu` keyword optionally specifies the MTU value of the VPLS or VPWS transport payload packet (customer packet). The MTU value is exchanged with VPLS-configured peer nodes. All VPLS peer nodes must be configured with the same MTU value. If the MTU values do not match, PWs cannot be established between peers. The MTU values are signaled during PW establishment so that endpoints can verify that MTU settings are equivalent before establishing the PW. By default the MTU is set to 1500. The configurable MTU range is 1492 through 9216. Changing the MTU setting causes established PWs to terminate. Payload packets might be dropped if the VPLS or VPWS MTU setting is greater than the *MPLS* MTU setting for the PW interface.

**Note**

The maximum MTU value supported depends on the current configuration options. For more information, see “Configuring the Layer 2 VPN MTU” in the [Switch Engine 32.3 User Guide](#).

The `l2vpn` keyword was introduced in ExtremeXOS Release 12.4 and is required when enabling a VPWS. For backward compatibility, the `l2vpn` keyword is optional when enabling a VPLS. However, this keyword will be required in a future release, so we recommend that you use this keyword for new configurations and scripts.

Example

The following commands change the various parameters of a particular VPLS:

```
configure vpls vpls1 dot1q ethertype 0x8508
configure vpls vpls1 dot1q ethertype 0x8509 mtu 2500
configure vpls vpls1 dot1q tag exclude mtu 2430
configure vpls vpls1 dot1q mtu 2500
```

History

This command was first available in ExtremeXOS 11.6.

The `l2vpn` and `vpws` keywords were first available in ExtremeXOS 12.4.

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure lacp member-port priority

```
configure lacp member-port port priority port_priority
```

Description

Configures the member port of an LACP to ensure the order that ports are added to the aggregator. The lower value you configure for the port's priority, the higher priority that port has to be added to the aggregator.

Syntax Description

<i>port</i>	Specifies the LACP member port that you are specifying the priority for.
<i>port_priority</i>	Specifies the priority you are applying to this member port to be assigned to the LACP aggregator. The range is from 0 to 65535; the default is 0. The lower configured value has higher priority to be added to the aggregator.

Default

The default priority is 0.

Usage Guidelines

The port must be added to the *LAG* prior to configuring it for LACP. The default value is 0, or highest priority.

You can configure the port priority to ensure the order in which LAG ports join the aggregator. If you do not configure this parameter, the lowest numbered ports in the LAG are the first to be added to the aggregator; if there are additional ports configured for that LAG, they are put in standby mode.

Use this command to override the default behavior and ensure the order in which LAG ports are selected. Also, if more than one port is configured with the same priority, the lowest numbered port joins the aggregator.

Example

The following command sets the port priority for the LAG port 5:1 to be 55 (which will probably put that port in standby initially):

```
configure lacp member-port 5:1 priority 55
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ldap domain

```
configure ldap domain domain_name [default | non-default]
```

Description

This command is used to configure a previously added LDAP domain as default or non-default. If a domain is configured as default, older default domain, if any, will no longer be default since once only one domain can be default at a time.

Syntax Description

<i>domain_name</i>	Name of domain to be configured.
--------------------	----------------------------------

Default

N/A.

Usage Guidelines

Use this command to configure an LDAP domain as default or non-default.

Example

This command marks the LDAP domain sales.XYZCorp.com as the default domain.

```
configure ldap domain sales.XYZCorp.com default
```

History

This command was first available in ExtremeXOS 15.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ldap domain add server

```
configure ldap {domain domain_name} add server [host_ipaddr | host_name]  
           {server_port} {client-ip client_ipaddr} {vr vr_name} {encrypted sasl  
           digest-md5}
```

Description

This command adds an LDAP server under an LDAP domain and configures the parameters for contacting the server.

Syntax Description

<i>domain_name</i>	Specifies the LDAP domain under which this server should be added.
<i>host_ipaddr</i>	Specifies a IP address for an LDAP server to add.
<i>host_name</i>	Specifies a DNS hostname for an LDAP server to add.
<i>server_port</i>	Specifies a port number for the LDAP service. The default port number is 389.
<i>client_ipaddr</i>	Specifies the LDAP client IP address, which should be set to the IP address of the interface that will connect to the LDAP server.
<i>vr_name</i>	Specifies the VR name for the interface that will connect to the LDAP server. The default VR for LDAP client connections is <i>VR-Mgmt</i> .
encrypted sasl digest-md5	<p>Specifies that the LDAP client uses Digest RSA Data Security, Inc. MD5 Message-Digest Algorithm encryption over SASL (Simple Authentication and Security Layer) to communicate with the LDAP server. Note that this mechanism encrypts only the password credentials, and the LDAP information exchange uses plain text.</p> <p>Note: To support Digest RSA Data Security, Inc. MD5 Message-Digest Algorithm over SASL, the LDAP client (bind user) password must be stored using 'reverse encryption,' and the host_name should be configured as the fully-qualified host name for the LDAP server.</p>

Default

client-ipaddr is optional. If client-ipaddr is not specified, the LDAP client looks up the interface through which the LDAP server can be reached.

If vr_name is not specified, the LDAP client assumes it to be VR-Mgmt.

If "encrypted sasl digest-md5" is not specified, the LDAP client talks to the LDAP server using plain text.

Usage Guidelines

You can configure up to 8 LDAP servers under one LDAP domain. The LDAP servers are contacted in the order of configuration. If the first server does not respond before the timeout period expires, the second server is contacted. This process continues until an LDAP server responds, and then the responding server marked as 'active'. Subsequent LDAP requests for that LDAP domain are sent to the 'active' server.



Note

If the switch cannot resolve the host name using a DNS server, the switch rejects the command and generates an error message.

As of 15.2, the "identity-management" keyword is now optional in this command.

Example

The following command configures LDAP client access to LDAP server LDAP1 using encrypted authentication:

```
* Switch.6 # configure ldap add server LDAP1 client-ip 10.10.2.1
encrypted sasl digest-md5
```

The following command adds the LDAP server LDAPServer1.sales.XYZCorp.com under the domain sales.XYZCorp.com and configures the LDAP client to contact it over VR-Default. It also configures the LDAP client to communicate with the server using digest-md5 encryption over SASL.

```
configure ldap domain sales.XYZCorp.com add server LDAPServer1.sales.XYZCorp.com vr VR-
Default encrypted sasl digest-md5
```

The following command adds the LDAP server 192.168.1.1 under the domain sales.XYZCorp.com and also configures the LDAP client to contact it through the interface 10.10.10.1 over VR-Mgmt.

```
configure ldap domain sales.XYZCorp.com add server 192.168.1.1 client-ip 10.10.10.1
```

History

This command was first available in ExtremeXOS 12.5.

This command was modified in ExtremeXOS 15.2 to make the identity management keyword optional.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ldap domain base-dn

```
configure ldap {domain [domain_name | all]} base-dn [base_dn | none |
default
```

Description

Configures the LDAP base-dn to be used while searching an user under an LDAP domain.

Syntax Description

<i>domain_name</i>	Specifies the LDAP domain for which this base-dn is to be configured.
<i>base_dn</i>	Specifies the LDAP base domain under which the users are to be searched.
none	Specifies the LDAP root domain as the location under which the users are to be searched.
default	Restores the base_dn to it default value i.e., same as the domain name.

Default

By default base-dn is assumed to be the same as the domain name unless configured otherwise.

If a domain is not specified, the base-dn is configured for the default domain.

Usage Guidelines

LDAP base-dn is the LDAP directory root under which the users are to be searched. By default base-dn is assumed to be the same as the domain name.

For users upgrading from ExtremeXOS 15.1 and older versions, a domain is created with the same name as the base-dn in the older configuration. This domain is marked as the default domain. This can be changed later if required.

Example

The following commands configure the base-dn for the domain sales.XYZCorp.com.

The base-dn configured as XYZCorp.com means that XYZCorp.com is the base location to search for user information.

```
* Switch.11 # configure ldap domain sales.XYZCorp.com base-dn XYZCorp.com
```

The base-dn configured as none means that the directory root is the base location to search for user information.

```
* Switch.12 # configure ldap domain sales.XYZCorp.com base-dn none
```

History

This command was first available in ExtremeXOS 12.5.

This command was modified in ExtremeXOS 15.2 to add the **{domain [domain_name | all]}** option.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ldap domain bind-user

```
configure ldap {domain [domain_name | all]} bind-user [user_name  
               {encrypted encrypted_password} | password | anonymous]
```

Description

Configures the LDAP client credentials required for the switch to access an LDAP server.

Syntax Description

<i>domain_name</i>	Specifies the LDAP domain for which this bind-user is to be configured.
<i>user_name</i>	Specifies the user name for LDAP server access.
encrypted	Indicates that the specified password is encrypted.
<i>password</i>	Specifies the user password for LDAP server access. Note: To support Digest RSA Data Security, Inc. <i>MD5</i> Message-Digest Algorithm over SASL, the password must be stored using 'reverse encryption.'
anonymous	Specifies user anonymous for LDAP server access.

Default

If no domain is specified, the bind-user is configured for the default domain.

Usage Guidelines

The bind-user is an LDAP user who has read access to user information in the LDAP directory.

On many newer directory servers "anonymous" access is disabled. You may also find that though the LDAP bind succeeds, the anonymous user might be denied read access to user information.

Example

The following command configures the LDAP bind user as jsmith with password Extreme for the domain sales.XYZCorp.com:

```
* Switch.14 # configure ldap domain sales.XYZCorp.com bind-user jsmith password Extreme
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ldap domain delete server

```
configure ldap {domain [domain_name |all]} delete server [host_ipaddr |
    host_name] {server_port} {vr vr_name}
```

Description

This command is used to delete one or all LDAP servers from one or all LDAP domains.

Syntax Description

<i>domain_name</i>	Specifies the LDAP domain from which this server is to be deleted.
all	Specifies that all configured LDAP servers are to be deleted.
<i>host_ipaddr</i>	Specifies the IP address of the LDAP server to delete.
<i>host_name</i>	Specifies a DNS hostname of the LDAP server to delete.
<i>server_port</i>	Specifies a port number for the LDAP service to delete. The default port number is 389.
vr <i>vr_name</i>	Specifies the virtual router to delete.

Default

If a domain is not specified, the server(s) under default domain is deleted.

Usage Guidelines

None.

Example

The following command deletes the LDAP server LDAPServer1.sales.XYZCorp.com from the domain sales.XYZCorp.com:

```
* Switch.8 # configure ldap domain sales.XYZCorp.com delete server
LDAPServer1.sales.XYZCorp.com
```

The following command deletes all LDAP servers from all LDAP domains:

```
* Switch.8 # configure ldap domain all delete server all
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ldap domain netlogin

```
configure ldap { domain [ domain_name | all ] } [enable|disable]
netlogin [dot1x | mac | web-based]
```

Description

Enables or disables LDAP queries for the specified type of network login users.

Syntax Description

<i>domain_name</i>	Specifies the LDAP domain for which this configuraton is to be applied.
dot1x	Enables or disables LDAP queries for dot1x network login.
mac	Enables or disables LDAP queries for MAC network login.
web-based	Enables or disables LDAP queries for web-based network login.

Default

LDAP queries are enabled for all types of network login.

Usage Guidelines

It may be necessary to disable LDAP queries for specific type of netlogin user, for example, netlogin mac users, whose username is the same as mac address. The LDAP directory might not contain useful information about these type of users and unnecessary LDAP queries can be avoided.



Note

LDAP queries are not sent for locally authenticated network login users.

Example

The following command enables LDAP queries for MAC network login:

```
* Switch.99 # configure ldap enable netlogin mac
```

The following command disables LDAP queries for dot1x network login:

```
* Switch.99 # configure ldap disable netlogin dot1x
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ldap hierarchical-search-oid

```
configure ldap {domain [domain_name|all]} hierarchical-search-oid [ldap-  
matching-rule-in-chain | oid | none]
```

Description

Configures an OID to perform a hierarchical search if the LDAP server requires it.

Syntax Description

<i>domain_name</i>	Domain name on which to configure ldap.
all	All domains.
ldap-matching-rule-in-chain	Configures the OID 1.2.840.113556.1.4.1941.
<i>oid</i>	Object identifier.
none	Specifies that LDAP query should not include any OID for hierarchical search.

Default

N/A.

Usage Guidelines

Use this command to configure an OID to perform a hierarchical search if the LDAP requires it. The OID supplied with this command will be used to form the LDAP query. If a server does not require extended control OID, the none option can be selected.

Example

```
configure ldap domain abc.com hierarchical-search-oid ldap_matching_rule_in_chain
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure lldp management-address

```
configure lldp management-address [[[{vlan} vlan_name | vlan vlan_id]  
{primary-ip | secondary-ip secondary_ip_address}] | mac-address]
```

Description

Configures a specified VLAN's IP address as the management address to be advertised by [LLDP](#).

Syntax Description

vlan	Specifies a VLAN for the management IP address.
<i>vlan_name</i>	Specifies a VLAN name for the management IP address (default is "Mgmt").
<i>vlan_id</i>	Specifies a VLAN ID for the management IP address.
primary-ip	LLDP advertises the primary IP address of the specified VLAN (default). The specified VLAN must be already configured with at-least one primary IPv4 address.
secondary-ip	Specifies that LLDP advertises the secondary IP address of the specified VLAN. The specified secondary IP address must already be configured on the specified VLAN. Note: LLDP does not recognize IPv6 addresses in this field.
<i>secondary_ip_address</i>	Specifies the secondary IP address of the specified VLAN.
mac-address	Specifies that LLDP advertises the MAC address of the switch. This is the default behavior if Management VLAN IP address is not configured and no VLAN is specified by this command.

Default

The system MAC address is advertised by default if the Management VLAN IP address is not configured and no VLAN is specified by this command.

By default, the Management VLAN's IP address is advertised by LLDP.

If you do not specify, LLDP advertises the primary IP address of the specified VLAN.

Usage Guidelines

If the Management VLAN IP address is not configured, LLDP and [CDP \(Cisco Discovery Protocol\)](#) advertise the system MAC address as the management address in their management TLV, which makes the network device not accessible. If the Management VLAN IP address is not configured, you can specify any user-defined VLAN's IP address or front panel port VLAN's IP address as the management address for LLDP and CDP protocols.

This command dictates the management address to be advertised by the LLDP protocol; the equivalent command for CDP is [configure cdp management-address](#) on page 340.

To use this command, the specified VLAN must already exist. The management IP address configuration is removed if the specified VLAN is deleted, or if the primary IP address of the specified VLAN is deleted (if **primary-ip** configured), or if the specified secondary IP address of the specified VLAN is deleted (if **secondary-ip** configured).

If **primary-ip** is configured and the specified VLAN has multiple primary IP addresses (IPv4 and IPv6), then LLDP advertises the first primary IP address that exists in the address table.

If **secondary-ip** is configured and the specified VLAN has multiple secondary IP addresses, then LLDP advertises only the specified secondary IP address of the configuration.



Note
LLDP does not recognize IPv6 addresses in this field.

Example

The following example configures the primary IP address of the VLAN "vlan1" as the management address to be advertised by LLDP protocol:

```
configure lldp management-address vlan vlan1 primary-ip
```

History

This command was first available in ExtremeXOS 22.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure lldp med fast-start repeat-count

```
configure lldp med fast-start repeat-count count
```

Description

The fast-start feature is automatically enabled when you enable the LLDP MED capabilities TLV. This command configures how many times, from 1 to 10, the switch sends out an LLDP MED packet with an interval of 1 second.

Syntax Description

<i>count</i>	Specifies the number of times the switch transmits LLDP MED TLVs each second (once it detects a neighbor transmitting LLDP MED TLVs). The range is 1 to 10.
--------------	---

Default

3.

Usage Guidelines

When the switch detects a MED-capable device, this count determines how many times the switch sends a LLDP MED TLVs with an interval of 1 second. The fast-start feature enables the MED-capable

device to quickly learn information; this command changes the value from the default 3. The fast-start feature is automatically enabled when you enable the LLDP MED capabilities TLV.



Note

After you configure the LLDP MED capability TLV, the fast-start feature automatically runs. To configure the LLDP MED capability TLV, use the `configure lldp ports [all | port_list] [advertise | no-advertise] vendor-specific med capabilities` command.

Example

The following command configures fast learning on the switch to a value of 2:

```
configure lldp med fast-start repeat-count 2
```

History

This command was first available in ExtremeXOS 11.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure lldp ports dcbx add application

```
configure lldp ports [all | port_list] dcbx add application [name
    application_name | ethertype ethertype_value | L4-port port_number |
    tcp-port port_number | udp-port port_number] priority priority_value
```

Description

Configures an application priority to be advertised to DCBX end stations.

Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.
<i>application_name</i>	Specifies an application. Supported values are: <ul style="list-style-type: none"> fcoe—Fiber Channel Over Ethernet (FCoE). fip—FCoE Initiation Protocol (FIP). iscsi—Internet Small Computer System Interface (iSCSI).
<i>ethertype_value</i>	Specifies an ethertype value in the range of 1536 to 65535.
L4-port <i>port_number</i>	Specifies a Layer 4 port number in the range of 0 to 65535. Supported Layer4 protocols include TCP, SCTP, UDP, and DCCP.
tcp-port <i>port_number</i>	Specifies a TCP port number in the range of 0 to 65535.

udp-port <i>port_number</i>	Specifies a UDP port number in the range of 0 to 65535.
priority_value	Specifies a priority in the range of 0 to 7.

Default

N/A.

Usage Guidelines

This command configures the switch to advertise the priority that an end station should use for the specified application or port number. The priority number is mapped to an 802.1p value, which determines how the switch manages traffic from that application or port.

The switch supports a maximum of 8 DCBX applications per port. If an application configuration already exists on the specified port or ports, the priority is updated to the new value. If the maximum number of applications for a port is exceeded, the switch logs an error message.

Example

The following command configures the switch to advertise priority 4 for the iSCSI application on ports 1 to 24:

```
configure lldp ports 1-24 dcbx add application name iscsi priority 4
```

The following command configures the switch to advertise priority 3 for ethertype value 34525 on port 1:

```
configure lldp ports 1 dcbx add application ethertype 34525 priority 3
```

The following command configures the switch to advertise priority 6 for Layer 4 port 992 on port 1:

```
configure lldp ports 1 dcbx add application L4-port 992 priority 6
```

History

This command was first available in ExtremeXOS 12.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure lldp ports dcbx delete application

```
configure lldp ports [all | port_list] dcbx delete application [all-applications | name application_name | ethertype ethertype_value | L4-port port_number | tcp-port port_number | udp-port port_number]
```

Description

Removes the priority configuration for one or all applications from the specified ports.

Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.
<i>application_name</i>	Specifies an application. Supported values are: <ul style="list-style-type: none"> fcoe—Fiber Channel Over Ethernet (FCoE). fip—FCoE Initiation Protocol (FIP). iscsi—Internet Small Computer System Interface (iSCSI).
<i>ethertype_value</i>	Specifies an ethertype value in the range of 1536 to 65535.
L4-port <i>port_number</i>	Specifies a Layer 4 port number in the range of 0 to 65535. Supported Layer4 protocols include TCP, SCTP, UDP, and DCCP.
tcp-port <i>port_number</i>	Specifies a TCP port number in the range of 0 to 65535.
udp-port <i>port_number</i>	Specifies a UDP port number in the range of 0 to 65535.

Default

N/A.

Usage Guidelines

This command configures the switch to advertise the priority that an end station should use for the specified application or port number. The priority number is mapped to an 802.1p value, which determines how the switch manages traffic from that application or port.

If an application configuration already exists on the specified port or ports, the priority is updated to the new value.

Example

The following command removes the priority configuration for Layer 4 port 30 on port 23:

```
configure lldp ports 23 dcbx delete application L4-port 30
```

History

This command was first available in ExtremeXOS 12.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure lldp ports management-address

```
configure lldp ports [all | port_list] [advertise | no-advertise]
management-address
```

Description

Configures the LLDP port to advertise or not to advertise management address information to its neighbors.

Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.

Default

No advertise.

Usage Guidelines

With ExtremeXOS, you can only add one management address TLV per LLDPDU and the information must be the IP address configured on the management VLAN. If no IP address is assigned to the management VLAN, the system sends the system MAC address. LLDP does not send out IPv6 addresses in this field.

Example

The following command advertises the management address information for port 1:5:

```
configure lldp ports 1:5 advertise management-address
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure lldp ports port-description

```
configure lldp ports [all | port_list] [advertise | no-advertise] port-
description
```

Description

Configures the LLDP port to advertise or not advertise port display information to its neighbors.

Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.

Default

No advertise.

Usage Guidelines

N/A.

Example

The following command configures port 1:7 to not advertise the port display information to neighbors:

```
configure lldp ports 1:7 no-advertise port-description
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure lldp ports system-capabilities

```
configure lldp ports [all | port_list] [advertise | no-advertise]  
    system-capabilities
```

Description

Configures the LLDP port to advertise or not to advertise its system capabilities to its neighbors.

Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.

advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.

Default

No advertise.

Usage Guidelines

When at least one VLAN exists with more than two ports, bridging is sent to enabled.

When at least one VLAN on the switch has IP forwarding enabled, the system automatically sets the router bit.

Example

The following command configures all ports to advertise system capability information to neighbors:

```
configure lldp ports all advertise system-capabilities
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure lldp ports system-description

```
configure lldp ports [all | port_list] [advertise | no-advertise]  
    system-description
```

Description

Configures the LLDP port to advertise or not to advertise its system description to its neighbors.

Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.

Default

Advertise.

Usage Guidelines

Although not mandatory according to the standard, this TLV is included in the LLDPDU by default when you enable LLDP.

When enabled, the system sends the following image (from the show version command) in the system description TLV:

```
ExtremeXOS version 11.2.0.12 v1120b12 by release-manager  
on Fri Mar 18 16:01:08 PST 2005
```

Example

The following command configures port 1:4 through port 1:8 to not advertise the system description information to neighbors:

```
configure lldp ports 1:4 - 1:8 no-advertise system-description
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure lldp ports system-name

```
configure lldp ports [all | port_list] [advertise | no-advertise]  
                  system-name
```

Default

Configures the LLDP port to advertise or not to advertise its system name to its neighbors.

Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.

Default

No advertise.

Usage Guidelines

N/A.

Example

The following command configures port 1:6 to advertise the system name to neighbors:

```
configure lldp ports 1:4 - 1:8 advertise system-name
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure lldp ports vendor-specific avaya-extreme call-server

The Avaya phone uses this proprietary LLDP TLV to learn the IP address(es) of the call server(s) to use.

```
configure lldp ports [all | port_list] [advertise | no-advertise]  
vendor-specific avaya-extreme call-server ip_address_1 {ip_address_2  
{ip_address_3 {ip_address_4 {ip_address_5 {ip_address_6 {ip_address_7  
{ip_address_8}}}}}}}}
```

Description

Configures the LLDP port to advertise or not advertise up to 8 call server IP addresses to its neighbors.

Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.
<i>ip_address_1... .8</i>	Specifies IP address of up to 8 call servers. Note: NOTE: This parameter does not apply when you configure the no-advertise parameter.

Default

No advertise.

Usage Guidelines

The Avaya phone uses this proprietary LLDP TLV for addressing information. You can configure the IP address for up to 8 call servers in a single TLV.

Example

The following command configures ports 1-5 to advertise two call server IP addresses to neighbors:

```
configure lldp ports 1-5 advertise vendor-specific avaya-extreme call-server 10.10.10.10
10.11.10.10
```

History

This command was first available in ExtremeXOS 11.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure lldp ports vendor-specific avaya-extreme dot1q-framing

```
configure lldp ports [all | port_list] [advertise | no-advertise]
vendor-specific avaya-extreme dot1q-framing [tagged | untagged |
auto]
```

Description

Configures the LLDP port to advertise or not advertise the 802.1q framing configuration to its neighbors. The Avaya phone uses this proprietary LLDP TLV information. In addition to this LLDP TLV, you must enable LLLDP as well as configure both the LLDP MED capabilities TLV and the LLDP network policy TLV.

Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.
tagged	Specifies to use tagging. NOTE: This parameter applies only when you use the advertise parameter.

untagged	Specifies not to use tagging. NOTE: This parameter applies only when you use the advertise parameter.
auto	Specifies following a predetermined sequence (see Usage Guidelines below). NOTE: This parameter applies only when you use the advertise parameter.

Default

No advertise.

Usage Guidelines

Before configuring this LLDP TLV, you must take the following steps:

- Enable LLDP using the `enable lldp ports` command.
- Enable the LLDP MED capabilities TLV using the `configure lldp ports vendor-specific med capabilities` command.
- Enable the LLDP MED network policy TLV using the `configure lldp ports vendor-specific med policy application` command.

This TLV is used to exchange information about Layer 2 priority tagging between the network connectivity device (switch) and the Avaya phone.

If you configure the TLV to advertise tagging, the phone uses tagging information, which it retrieves from the `configure lldp ports vendor-specific med policy application` command. If you configure the TLV to advertise untagged, the phone does not use any tagging, including 802.1q priority tagging.

If you configure the TLV to advertise auto, the phone cycles through the following sequence until an action is successful:

- Uses the configuration advertised by the LLDP MED network policy TLV, as configured by the `configure lldp ports vendor-specific med policy application` command.
- Uses the priority tagged frames configured by the phone's server.
- Sends the traffic untagged.

Example

The following command configures all ports to advertise the dot1q framing as untagged to neighbors:

```
configure lldp ports all advertise vendor-specific avaya-extreme dot1q-framing untagged
```

History

This command was first available in ExtremeXOS 11.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure lldp ports vendor-specific avaya-extreme file-server

```
configure lldp ports [all | port_list] [advertise | no-advertise]  
  vendor-specific avaya-extreme file-server ip_address_1 {ip_address_2  
    {ip_address_3 {ip_address_4}}}
```

Description

Configures the *LLDP* port to advertise or not advertise up to 4 file server IP addresses to its neighbors. The Avaya phone uses this proprietary LLDP TLV to learn the IP address(es) of the file server(s) to use.

Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.
<i>ip_address_1...4</i>	Specifies IP address of up to 4 file servers. NOTE: This parameter does not apply when you configure the no-advertise parameter.

Default

No advertise.

Usage Guidelines

The Avaya phone uses this proprietary LLDP TLV for addressing information. You can configure the IP address for up to 4 file servers in a single TLV.

Example

The following command configures all ports to advertise two file server IP addresses to neighbors:

```
configure lldp ports 1-5 advertise vendor-specific avaya-extreme call-server 10.20.10.10  
10.12.10.10
```

History

This command was first available in ExtremeXOS 11.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure lldp ports vendor-specific avaya-extreme poe-conservation-request

```
configure lldp ports [all | port_list] [advertise | no-advertise]  
vendor-specific avaya-extreme poe-conservation-request
```

Description

Configures the LLDP port to advertise or not advertise a requested conservation level. By default, the requested conservation value on this proprietary LLDP TLV is 0, which is no power conservation. This LLDP TLV is sent out only on PoE-capable Ethernet ports.

Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.

Default

No advertise.

Usage Guidelines

The switch sends this proprietary LLDP TLV to request a PD to go into a certain power conservation level or request the PD to go to the maximum conservation level. This LLDP TLV is transmitted only on PoE-capable ports.

When configured to advertise, the switch sends this TLV with a requested conservation power level of 0, which requests no power conservation. To temporarily change this conservation level, use the SNMP `lldpXAvExLocPortXPoEPSEPortReqLevel` object to set a new value; the reconfigured value is not saved over a reboot. (This SNMP object can be set from 0 to 243 or 255.)

Example

The following command configures all ports to advertise the currently requested conservation level to neighbors:

```
configure lldp ports all advertise vendor-specific avaya-extreme poe-conservation-request
```

History

This command was first available in ExtremeXOS 11.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure lldp ports vendor-specific dcbx

```
configure lldp ports [all | port_list] [advertise | no-advertise]  
    vendor-specific dcbx {ieee|baseline}
```

Description

Configures the LLDP port to advertise or not to advertise Data Center Bridging Exchange (DCBX) information to its neighbors.

Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.
ieee	Specifies the DCBX protocol defined in IEEE 802.1Qaz.
baseline	Specifies the DCBX protocol known as Baseline Version 1.01, which was defined before IEEE 802.1Qaz.

Default

No advertisement for both DCBX protocols.

Usage Guidelines

If you do not specify a protocol with this command, the advertise option enables advertisement for the IEEE 802.1Qaz protocol, and the no-advertise option disables advertisement for both protocols.

Example

The following command advertises DCBX information according to IEEE 802.1Qaz for port 1:5:

```
configure lldp ports 1:5 advertise vendor specific dcbx
```

The following command advertises DCBX information according to Baseline Version 1.01 for port 2:1:

```
configure lldp ports 2:1 advertise vendor specific dcbx baseline
```

The following command disables advertisement of DCBX information on all ports:

```
configure lldp ports all no-advertise vendor specific dcbx
```

History

This command was first available in ExtremeXOS 12.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure lldp ports vendor-specific dot1 port-protocol-vlan-ID

```
configure lldp ports [all | port_list] [advertise | no-advertise]
vendor-specific dot1 port-protocol-vlan-ID {vlan [all | vlan_name] }
```

Description

Configures the LLDP port to advertise or not advertise port VLAN information to its neighbors.

Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.
all	Specifies all <u>VLANs</u> on the port.
<i>vlan_name</i>	Specifies the VLAN on the port that you want to advertise.

Default

No advertise.

Usage Guidelines

When configured to advertise, the switch inserts a port and protocol VLAN ID TLV for each VLAN configured on the ports. The port and protocol VLAN ID TLV allows the port to advertise if it supports protocol and/or tagged VLANs, along with the associated tagged values. A separate TLV is sent for each VLAN that you want to advertise.

By default, once you configure this TLV, the system sends all protocol-based VLANs on the port. However, the LLDPDU cannot exceed 1500 bytes, so you should configure the port to advertise only the specified VLANs.



Note

The total LLDPDU size is 1500 bytes; any TLVs after that limit are dropped.

This TLV does not send information on the type of protocol that the VLAN has enabled; it just says whether the port is enabled or disabled for protocol-based VLANs. As Extreme Networks devices are always capable of supporting protocol-based VLANs, once you configure this TLV, the system always advertises support for these VLANs.

Example

The following command configures all ports to advertise port and protocol VLAN information to neighbors for all VLANs on all ports:

```
configure lldp ports all advertise vendor-specific dot1 port-protocol-vlan-id
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure lldp ports vendor-specific dot1 port-vlan-ID

```
configure lldp ports [all | port_list] [advertise | no-advertise]  
    vendor-specific dot1 port-vlan-ID
```

Description

Configures the LLDP port to advertise or not advertise port vlan ID information to its neighbors. This allows a VLAN bridge port to advertise the port VLAN identifier that is associated with untagged or priority-tagged frames.

Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.

Default

No advertise.

Usage Guidelines

The port VLAN ID TLV allows the port to transmit the VLAN ID associated with untagged VLANs. There can be only one port VLAN ID in each LLDPDU.

If no untagged VLANs are configured on the specified port, the TLV is not added to the LLDPDU, even if you configured this to advertise.

Example

The following command configures all ports to advertise port vlan ID information to neighbors:

```
configure lldp ports all advertise vendor-specific dot1 port-vlan-ID
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure lldp ports vendor-specific dot1 vlan-name

```
configure lldp ports [all | port_list] [advertise | no-advertise]  
    vendor-specific dot1 vlan-name {vlan [all | vlan_name]}
```

Description

Configures the *LLDP* port to advertise or not advertise *VLAN* name information to its neighbors. Use this TLV to advertise information for the tagged VLANs you want to specify on the port. This allows an IEEE 802.1Q-compatible 802 LAN station to advertise the assigned name of any VLAN with which it is configured.

Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.
vlan	Specifies all VLANs on the port.
<i>vlan_name</i>	Specifies the VLAN on the port that you want to advertise.

Default

No advertise.

Usage Guidelines

The VLAN name TLV sends the VLAN name and the tag used; it associates a name to a tag for the specified VLAN. This allows an IEEE 802.1Q-compatible 802 LAN station to advertise the assigned name of any VLAN with which it is configured.

You can enable this TLV for tagged and untagged VLANs. When you enable this TLV for tagged VLANs, the TLV advertises the IEEE 802.1Q tag for that VLAN. (For untagged VLANs, the internal tag is advertised.) You can specify exactly which VLANs to advertise.

When configured to advertise, the switch inserts a VLAN name TLV for every VLAN configured on the ports. By default, once you configure this TLV, the system sends all VLAN names on the port. However, each VLAN name can require up to 32 bytes and the LLDPDU cannot exceed 1500 bytes, so you should configure the port to advertise only the specified VLANs, using the keyword `vlan_name`.



Note

The total LLDPDU size is 1500 bytes; any TLVs after that limit are dropped.

Example

The following command configures all ports to not advertise VLAN name information to neighbors:

```
configure lldp ports all no-advertise vendor-specific dot1 vlan-name
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure lldp ports vendor-specific dot3 link-aggregation

```
configure lldp ports [all | port_list] [advertise | no-advertise]  
                vendor-specific dot3 link-aggregation
```

Description

Configures the LLDP port to advertise or not advertise link-aggregation capabilities to its neighbors.

Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.

Default

No advertise.

Usage Guidelines

When configured, this TLV is added to each LLDP port LLDPDU indicating the link-aggregation capabilities, status, and value of the master port of the load-sharing group.

Example

The following command configures port 1:12 to not advertise link-aggregation capabilities to neighbors:

```
configure lldp ports 1:12 no-advertise vendor-specific dot3 link-aggregation
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure lldp ports vendor-specific dot3 mac-phy

```
configure lldp ports [all | port_list] [advertise | no-advertise]  
                vendor-specific dot3 mac-phy
```

Description

Configures the LLDP port to advertise or not advertise MAC and physical layer capabilities to its neighbors. The capabilities include duplex and bit rate.

Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.

Default

No advertise.

Usage Guidelines

When configured, the system advertises information about the speed capabilities, as well as autonegotiation support and status, of the LLDP port.

Example

The following command configures all ports to advertise MAC/PHY capabilities to neighbors:

```
configure lldp ports all advertise vendor-specific dot3 mac-phy
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure lldp ports vendor-specific dot3 max-frame-size

```
configure lldp ports [all | port_list] [advertise | no-advertise]  
                vendor-specific dot3 max-frame-size
```

Description

Configures the LLDP port to advertise or not advertise its maximum frame size to its neighbors.

Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.

Default

No advertise.

Usage Guidelines

When jumbo frames are not enabled on the specified port, the TLV reports a value of 1518 once you configure it to advertise. If jumbo frames are enabled, the TLV inserts the configured value for the jumbo frames.

Example

The following command configures ports 1:12 and 1:13 to advertise the maximum frame size to neighbors:

```
configure lldp ports 1:12 - 1:13 advertise vendor-specific dot3 max-frame-size
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure lldp ports vendor-specific dot3 power-via-mdi

```
configure lldp ports [all | port_list] [advertise | no-advertise]  
    vendor-specific dot3 power-via-mdi {with-classification}
```

Description

Configures the LLDP port to advertise or not advertise Power over Ethernet (PoE) capabilities to its neighbors.

Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.
with-classification	Specifies to use LLDP for Data Link Layer Classification. This option is available only on PoE+ ports.

Default

No advertise.

Usage Guidelines

When configured, the system includes this TLV. We recommend enabling this TLV only on PoE-capable ports.

The following information is transmitted for LLDP ports with this TLV:

- Support PoE or not
- Port class
 - Power sourcing equipment (PSE)
 - Powered device (PD)
- Power pairs used to supply power
 - Signal
 - Spare
- Power status

- Support pairs control or not
- Power class
 - Class0
 - Class1
 - Class2
 - Class2
 - Class3
 - Class4

Data link layer classification allows fine-grained dynamic re-allocation of power based on changing needs. This feature is enabled by enabling LLDP (transmit and receive) and configuring transmission of the power-via-MDI TLV. The ExtremeXOS software sends an LLDPDU containing a power-via-MDI TLV within 10 seconds of DLL classification being enabled. A PD may request a new power value using an LLDPDU. The allocated power might be changed if a request is received and approved after a power review. The software responds with an allocated power value within 10 seconds of receipt of an LLDPDU with a different requested power from a PD. Power allocation can be controlled to a granularity of 0.1 watts. When DLL classification is enabled, it takes precedence over physical classification.



Note

For more information on advertising power support, see the [configure lldp ports vendor-specific med power-via-mdi](#) command.

Example

The following command configures all ports to advertise power capabilities to neighbors:

```
configure lldp ports all advertise vendor-specific dot3 power-via-mdi
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure lldp ports vendor-specific med capabilities

```
configure lldp ports [all | port_list] [advertise | no-advertise]  
vendor-specific med capabilities
```

Description

Configures the LLDP port to advertise or not advertise MED capabilities. This TLV must be enabled before any of the other MED TLVs can be enabled. Also, this TLV must be set to no-advertise after all other MED TLVs are set to no-advertise.

Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.

Default

No advertise.

Usage Guidelines

This command enables the LLDP media endpoint discovery (MED) capabilities TLV, which allows LLDP-MED network connectivity devices to definitively determine that particular endpoints support LLDP MED, and if so, to discover which LLDP MED TLVs the particular endpoint devices are capable of supporting and to which specific device class the device belongs to.

This TLV must be enabled before any of the other MED TLVs can be enabled; and this TLV must be set to no-advertise after all other MED TLVs are set to no-advertise.

As with all the LLDP MED TLVs, the switch sends this TLV only after it detects a MED-capable device on the port. The switch does not automatically send this TLV after it is enabled; the switch must first detect a MED-capable device on the port.



Note

Network connectivity devices wait to detect LLDP MED TLVs from endpoints before they send out LLDP MED TLVs; so L2 network connectivity devices do not exchange LLDP MED messages.

The following information is included in the LLDP MED capabilities TLV when it is transmitted:

- The supported LLDP MED TLVs—For Extreme Networks devices, these are capabilities, network policy, location, and extended power (extended power only advertised only on *PoE*-capable ports).
- The MED device type—For Extreme Networks devices, this is advertised as a network connectivity device (set to 4).

Example

The following command configures all ports to advertise MED capabilities to neighbors:

```
configure lldp ports all advertise vendor-specific med capabilities
```

History

This command was first available in ExtremeXOS 11.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure lldp ports vendor-specific med location-identification

```
configure lldp ports [all | port_list] [advertise | no-advertise]
vendor-specific med location-identification [coordinate-based
hex_value | civic-based hex_value | ecs-elin elin]
```

Description

Configures the LLDP port to advertise or not advertise MED location information. You configure up to 3 different location identifiers.

Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies to not send the information to neighbors.
coordinate-based	Specifies using the coordinate-based location identifier. This value is exactly 16 bytes long; see RFC 3825 for details.
<i>hex_value</i>	Enter a hexadecimal value with each byte separated by a colon. Or, you can obtain this value from a network management application. NOTE: This parameter is not used when the no-advertise parameter is configured.
civic-based	Specifies using the civic-based location identifier. This value must have a minimum length of 6 bytes; see RFC3825 for details.
ecs-elin	Specifies using the ecs location identifier. (Emergency Call Service, as defined in the TIA-TSB-146.)
<i>elin</i>	Enter a numerical string; the range is 10 to 25 characters. Or, you can obtain this value from a network management application. (See the TIA-TSB-146 standard for a definition of these numbers; also, the network management application must be able to handle the LLDP MED MIB.) NOTE: This parameter is not used when the no-advertise parameter is configured.

Default

No advertise.

Usage Guidelines

You might need to use a specific format for your specific VoIP implementation; see the VoIP manufacturer's manual for details.

You must configure the LLDP MED capabilities TLV before configuring this TLV. Configure the LLDP MED capabilities TLV using the `configure lldp ports [all | port_list] [advertise | no-advertise] vendor-specific med capabilities` command.

As with all the LLDP MED TLVs, the switch sends this TLV only after it detects a MED-capable device on the port. The switch does not automatically send this TLV after it is enabled; the switch must first detect a MED-capable device on the port.

Example

The following command configures all ports to advertise MED location information to neighbors using the ECS format:

```
configure lldp ports all advertise vendor-specific med location-identification ecs-elin
423233455676
```

History

This command was first available in ExtremeXOS 11.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure lldp ports vendor-specific med policy application

```
configure lldp ports [all | port_list] [advertise | no-advertise]
  vendor-specific med policy application [voice | voice-signaling
  | guest-voice | guest-voice-signaling | softphone-voice | video-
  conferencing | streaming-video | video-signaling] vlan vlan_name dscp
  dscp_value {priority-tagged}
```

Description

Configures the LLDP port to advertise or not advertise MED network policy TLVs. This TLV advertises VLAN configuration and associated Layer 2 and Layer 3 attributes that apply for a set of specific applications on that port. You can advertise up to 8 TLVs, each for a specific application, per port/VLAN. Each application type can exist only once per port. This TLV tells the endpoint the specific VLAN to use for the specific application, along with its unique priority.

Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.
advertise	Specifies to send the information to neighbors.

voice	Specifies voice application on specified port/VLAN(s).
voice-signaling	Specifies voice signaling application on specified port/VLAN(s).
guest-voice	Specifies guest voice application on specified port/VLAN(s).
guest-voice-signaling	Specifies guest voice signaling application on specified port/VLAN(s).
softphone-voice	Specifies soft phone voice application on specified port/VLAN(s).
video-conferencing	Specifies videoconferencing application on specified port/VLAN(s).
streaming-video	Specifies streaming video application on specified port/VLAN(s).
video-signaling	Specifies video signaling application on specified port/VLAN(s).
<i>vlan_name</i>	Specifies the VLAN the specified application is using. NOTE: This parameter does not apply when the no-advertise parameter is configured.
<i>dscp_value</i>	Specifies the DSCP value for the specified application. This is a 6-bit value from 0 to 63. NOTE: This parameter does not apply when the no-advertise parameter is configured.
priority-tagged	Use this if you want priority tagging, and the VLAN is configured as untagged on the port. (The endpoint sends out frames for the specified application with a tag of 0.) NOTE: This parameter does not apply when the no-advertise parameter is configured.

Default

No advertise.

Usage Guidelines

This command enables the LLDP MED network policy TLV, which allows network connectivity devices and endpoint devices to advertise VLAN configuration and associated Layer 2 and Layer 3 attributes that apply for a set of specific application on that port. This TLV can be enabled on a per port/VLAN basis. Each application type can exist only once on a port.

You can enable the transmission of a TLV policy for each application. A maximum of 8 TLVs can be enabled, and each can have a unique DSCP value and/or priority tagging.

You must configure the LLDP MED capabilities TLV before configuring this TLV. Configure the LLDP MED capabilities TLV using the `configure lldp ports [all | port_list] [advertise | no-advertise] vendor-specific med capabilities` command.

As with all the LLDP MED TLVs, the switch sends this TLV only after it detects a MED-capable device on the port. The switch does not automatically send this TLV after it is enabled; the switch must first detect a MED-capable device on the port.

The following information is transmitted for LLDP ports with this TLV:

- Application type

Used as configured.

- Unknown policy flag

Set to 0.

- Tagged flag

Set to tagged for tagged VLANs; set to untagged for untagged VLANs. By default, set to 0.

- VLAN ID

Copied from the VLAN. However, if you configure the priority-tagged parameter, this value is set to 0.

- Layer 2 priority

Copied from the VLAN priority.

- DSCP value

Uses the value configured in the dscp parameter.



Note

See the documentation provided by the manufacturer of connected devices regarding values.

Example

The following command configures all ports to advertise videoconferencing on the VLAN video with a DSCP of 7 to neighbors:

```
configure lldp ports all advertise vendor-specific med policy application video-  
conferencing vlan video dscp 7
```

History

This command was first available in ExtremeXOS 11.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure lldp ports vendor-specific med power-via-mdi

```
configure lldp ports [all | port_list] [advertise | no-advertise]  
vendor-specific med power-via-mdi
```

Description

Configures the *LLDP* port to advertise or not advertise MED power requirement details. This TLV can only be enabled on a *PoE*-capable port and is used for advanced power management between the MED network connectivity and endpoint devices.

Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.
advertise	Specifies to send the information to neighbors.
no-advertise	Specifies not to send the information to neighbors.

Default

No advertise.

Usage Guidelines

When enabled, this LLDP MED TLV advertises fine-grained power requirement details about PoE settings and support. This TLV can be enabled only on a PoE-capable port; the switch returns an error message if this TLV is configured for a non-PoE-capable port.

You must configure the LLDP MED capabilities TLV before configuring this TLV. Configure the LLDP MED capabilities TLV using the `configure lldp ports [all | port_list] [advertise | no-advertise] vendor-specific med capabilities` command.

As with all the LLDP MED TLVs, the switch sends this TLV only after it detects a MED-capable device on the port. The switch does not automatically send this TLV after it is enabled; the switch must first detect a MED-capable device on the port.



Note

For additional information on power support, see the `configure lldp ports vendor-specific dot3 power-via-mdi` command.

The following information is transmitted for LLDP MED PoE-capable ports with this TLV:

- Power type
 - Set to PSE.
- Power source
 - Set to primary power source.
- Power priority
 - Taken from PoE port configuration.
- Power value
 - Taken from PoE port configuration.

Example

The following command configures all ports to advertise MED power information to neighbors:

```
configure lldp ports all advertise vendor-specific med power-via-mdi
```

History

This command was first available in ExtremeXOS 11.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure lldp reinitialize-delay

```
configure lldp reinitialize-delay seconds
```

Description

Configures the delay before the receive state machine is reinstalled once the *LLDP* transmit mode has been disabled.

Syntax Description

<i>seconds</i>	Specifies the delay that applies to the reinitialization attempt. The range is 1 to 10 seconds.
----------------	---

Default

2 seconds.

Usage Guidelines

N/A.

Example

The following command configures a reinitialization delay of 10 seconds:

```
configure lldp reinitialize-delay 10
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure lldp snmp-notification-interval

```
configure lldp snmp-notification-interval seconds
```

Description

Configures the allowed interval at which *SNMP* notifications are sent.

Syntax Description

<i>seconds</i>	Specifies the interval at which <i>LLDP</i> SNMP notifications are sent. The range is 5 to 3600 seconds.
----------------	--

Default

5 seconds.

Usage Guidelines

This is a global timer. If one port sends a notification, no notifications for other ports go out for the configured interval.

Example

The following command configures an interval of 60 seconds for LLDP SNMP notifications:

```
configure lldp snmp-notification-interval 60
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure lldp transmit-delay

```
configure lldp transmit-delay [ auto | seconds]
```

Description

Configures the delay time between successive frame transmissions initiated by a value change or status change in any of the *LLDP* local systems Management Information Base (MIB).

The auto option uses a formula ($0.25 * \text{transmit-interval}$) to calculate the number of seconds.

Syntax Description

auto	Uses the formula ($0.25 * \text{transmit-interval}$) to calculate the seconds.
<i>seconds</i>	Specifies the interval at which LLDP notifications are sent. The range is 1 to 8291.

Default

2 seconds.

Usage Guidelines

This is the timer between triggered updates.

Example

The following command configures the delay between LLDP frame transmissions for triggered updates to be automatically calculated:

```
configure lldp transmit-delay auto
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure lldp transmit-hold

```
configure lldp transmit-hold hold
```

Description

Calculates the actual time-to-live (TTL) value used in the LLDPDU messages.

The formula is $\text{transmit-interval} * \text{transmit-hold}$; by default the TTL value is $(30 * 4)$ 120 seconds.

Syntax Description

<i>hold</i>	Used to calculate the TTL value; the range is 2 to 10.
-------------	--

Default

4.

Usage Guidelines

N/A.

Example

The following command configures the transmit-hold value (which is used to calculate the TTL of the *LLDP* packets) to 5:

```
configure lldp transmit-hold 5
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure lldp transmit-interval

```
configure lldp transmit-interval seconds
```

Description

Configures the periodic transmittal interval for LLDPDUs.

Syntax Description

<i>seconds</i>	Specifies the time between LLDPDU transmissions. The range is 5 to 32768.
----------------	---

Default

30 seconds.

Usage Guidelines

N/A.

Example

The following command configures a transmittal interval of 20 seconds for LLDPDUs.

```
configure lldp transmit-interval 20
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure log display

```
configure log display severity {only}
```

Description

Configures the real-time log-level message to display.

Syntax Description

<i>severity</i>	Specifies a message severity. Severities include critical, error, warning, notice, info, debug-summary, debug-verbose, and debug-data.
only	Specifies only log messages of the specified severity level.

Default

If not specified, messages of all severities are displayed on the console display.

Usage Guidelines

You must enable the log display before messages are displayed on 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.

Severity filters the log to display messages with the selected severity or higher (more critical). Severities include critical, error, warning, info, notice, debug-summary, debug-verbose, and debug-data.

You can also control log data to different targets. The command equivalent to [configure log display](#) is the following:

```
configure log target console-display severity severity
```

To display the current configuration of the log display, use the following command:

```
show log configuration target console-display
```

In a stack, this command is applicable only to Master and Backup nodes and not applicable to the standby nodes.

Example

The following command configures the system log to maintain a running real-time display of log messages of critical severity or higher:

```
configure log display critical
```

The following command configures the system log to maintain a running real-time display of only log messages of critical severity:

```
configure log display critical only
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure log filter events

```
configure log filter name [add | delete] {exclude} events [event-condition | [all | event-component] {severity severity {only}}]
```

Description

Configures a log filter to add or delete detailed feature messages based on a specified set of events.

In a stack, this command is applicable only to Master and Backup nodes and not applicable to the standby nodes.

Syntax Description

<i>name</i>	Specifies the filter to configure.
add	Add the specified events to the filter.
delete	Remove the specified events from the filter.
exclude	Events matching the specified events will be excluded.
<i>event-condition</i>	Specifies an individual event.
all	Specifies all components and subcomponents.
<i>event-component</i>	Specifies all the events associated with a particular component.

<i>severity</i>	Specifies the minimum severity level of events (if the keyword only is omitted).
only	Specifies only events of the specified severity level.

Default

If the exclude keyword is not used, the events will be included by the filter. If severity is not specified, then the filter will use the component default severity threshold (see the note [note: If no severity is specified when delete or exclude is specified, severity all is used](#) when delete or exclude is specified).

Usage Guidelines

This command controls the incidents that pass a filter by adding, or deleting, a specified set of events. If you want to configure a filter to include or exclude incidents based on event parameter values (for example, MAC address or *BGP* Neighbor) see the command [configure log filter events match](#).

When the add keyword is used, the specified event name is added to the beginning of the filter item list maintained for this filter. The new filter item either includes the events specified, or if the exclude keyword is present, excludes the events specified.

The delete keyword is used to remove events from the filter item list that were previously added using the add command. All filter items currently in the filter item list that are identical to, or a subset of, the set of events specified in the delete command will be removed.

Event Filtering Process

From a logical standpoint, the filter associated with each enabled log target is examined to determine whether a message should be logged to that particular target. The determination is made for a given filter by comparing the incident with the most recently configured filter item first. If the incident matches this filter item, the incident is either included or excluded, depending on whether the exclude keyword was used. Subsequent filter items on the list are compared if necessary. If the list of filter items has been exhausted with no match, the incident is excluded.

Events, Components, and Subcomponents

As mentioned, a single event can be included or excluded by specifying the event's name. Multiple events can be added or removed by specifying an ExtremeXOS component name plus an optional severity. Some components, such as BGP, contain subcomponents, such as Keepalive, which is specified as BGP.Keepalive. Either components or subcomponents can be specified. The keyword all in place of a component name can be used to indicate all ExtremeXOS components.

Severity Levels

When an individual event name is specified following the events keyword, no severity value is needed since each event has pre-assigned severity. When a component, subcomponent, or the all keyword is specified following the events keyword, a severity value is optional. If no severity is specified, the severity used for each applicable subcomponent is obtained from the pre-assigned severity threshold

levels for those subcomponents. For example, if *STP* were specified as the component, and no severity is specified for the add of an include item, then only messages with severity of error and greater would be passed, since the threshold severity for the STP component is error. If STP.InBPDU were specified as the component, and no severity is specified, then only messages with severity of warning and greater would be passed, since the threshold severity for the STP.InBPDU subcomponent is warning. Use the `show log components` command to see this information.

The severity keyword all can be used as a convenience when delete or exclude is specified. The use of delete (or exclude) with severity all deletes (or excludes) previously added events of the same component of all severity values.

**Note**

If no severity is specified when delete or exclude is specified, severity all is used.

If the only keyword is present following the severity value, then only the events in the specified component at that exact severity are included. Without the only keyword, events in the specified component at that severity or more urgent are included. For example, using the option severity warning implies critical, error, or warning events, whereas the option severity warning only implies warning events only. Severity all only is not a valid choice.

Any EMS events with severity debug-summary, debug-verbose, or debug-data will not be logged unless debug mode is enabled. See the command `enable log debug-mode`.

Filter Optimization

Each time a configure log filter command is issued for a given filter name, the events specified are compared against the current configuration of the filter to try to logically simplify the configuration.

For example, if the command:

```
configure log filter bgpFilter1 add events bgp.keepalive severity error only
```

were to be followed by the command:

```
configure log filter bgpFilter1 add events bgp severity info
```

the filter item in the first command is automatically deleted since all events in the BGP.Keepalive subcomponent at severity error would be also included as part of the second command, making the first command redundant.

More Information

See the command `show log` for more information about severity levels.

To get a listing of the components present in the system, use the following command:

```
show log components
```

To get a listing of event condition definitions, use the following command:

```
show log events
```

To see the current configuration of a filter, use the following command:

```
show log configuration filter {filter name}
```

Example

The following command adds all STP component events at severity info to the filter mySTPFilter:

```
configure log filter myStpFilter add events stp severity info
```

The following command adds the STP.OutBPDU subcomponent, at the pre-defined severity level for that component, to the filter myStpFilter:

```
configure log filter myStpFilter add events stp.outbpdu
```

The following command excludes one particular event, STP.InBPDU.Drop, from the filter:

```
configure log filter myStpFilter add exclude events stp.inbpdu.drop
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure log filter events match

```
configure log filter name [add | delete] {exclude} events [event-  
condition | [all | event-component] {severity severity {only}}]  
[match | strict-match] type value
```

Description

Configures a log filter to add or delete detailed feature messages based on a specified set of events and match parameter values.

In a stack, this command is applicable only to Master and Backup nodes and not applicable to the standby nodes.

Syntax Description

<i>name</i>	Specifies the filter to configure.
add	Add the specified events to the filter.
delete	Remove the specified events from the filter.
exclude	Events matching the filter will be excluded.
<i>event-condition</i>	Specifies the event condition.
all	Specifies all events.

<i>event-component</i>	Specifies all the events associated with a particular component.
<i>severity</i>	Specifies the minimum severity level of events (if the keyword only is omitted).
only	Specifies only events of the specified severity level.
match	Specifies events whose parameter values match the <i>type value</i> pair.
strict-match	Specifies events whose parameter values match the <i>type value</i> pair, and possess all the parameters specified.
<i>type</i>	Specifies the type of parameter to match. For more information about types and values see Types and Values .
<i>value</i>	Specifies the value of the parameter to match. For more information about types and values see Types and Values .

Default

If the **exclude** keyword is not used, the events will be included by the filter. If severity is not specified, then the filter will use the component default severity threshold (see the note on [note: If no severity is specified when delete or exclude is specified, severity all is used](#) when delete or exclude is specified).

Usage Guidelines

This command controls the incidents that pass a filter by adding or deleting a specified set of events that match a list of *type value* pairs. This command is an extension of the command [configure log filter events](#), and adds the ability to filter incidents based on matching specified event parameter values to the event.

See the [configure log filter events](#) command [configure log filter events](#) for more information on specifying and using filters, on event conditions and components, and on the details of the filtering process. The discussion here is about the concepts of matching *type value* pairs to more narrowly define filters.

Types and Values

Each event in ExtremeXOS is defined with a message format and zero or more parameter types. The [show log events](#) command [show log events](#) can be used to display event definitions (the event text and parameter types). The syntax for the parameter types (represented by *type* in the command syntax above) is:

```
[address-family [ipv4-multicast | ipv4-unicast | ipv6-multicast | ipv6-unicast] | bgp-neighbor ip address | bgp-routerid ip address | eaps eaps domain name | {destination | source} [ipaddress ip address | L4-port | mac-address ] | esrp esrp domain name | {egress | ingress} [slot slot number | port port_list] | ipaddress ip address | L4-port L4-port | mac-address mac_address | netmask netmask | number number | port port_list |
```

```
process process name | slot slotid | string exact string to be matched |
vlan vlan name | vlan tag vlan tag]
```

**Note**

The slot parameters are available only on SummitStacks.

Beginning with ExtremeXOS 11.2, you can specify the `ipaddress` type as IPv4 or IPv6, depending on the IP version. The following examples show how to configure IPv4 addresses and IPv6 addresses:

- IPv4 address.

To configure an IP address, with a mask of 32 assumed, use the following command:

```
configure log filter myFilter add events all match ipaddress 12.0.0.1
```

To configure a range of IP addresses with a mask of 8, use the following command:

```
configure log filter myFilter add events all match ipaddress 12.0.0.0/8
```

- IPv6 address.

To configure an IPv6 address, with a mask of 128 assumed, use the following command:

- `configure log filter myFilter add events all match ipaddress 3ffe::1`
- To configure a range of IPv6 addresses with a mask of 16, use the following command:
- `configure log filter myFilter add events all match ipaddress 3ffe::/16`

- IPv6 scoped address.

IPv6 scoped addresses consist of an IPv6 address and a VLAN. The following examples identify a link local IPv6 address.

To configure a scoped IPv6 address, with a mask of 128 assumed, use the following command:

**Note**

In the previous example, if you specify the VLAN name, it must be a full match; wild cards are not allowed.

The *value* depends on the parameter type specified. As an example, an event may contain a physical port number, a source MAC address, and a destination MAC address. To allow only those incidents with a specific source MAC address, use the following in the command:

```
configure log filter myFilter add events aaa.radius.requestInit security
notice match source mac-address 00:01:30:23:C1:00 configure log filter
myFilter add events bridge severity notice match source mac-address
00:01:30:23:C1:00
```

The string type is used to match a specific string value of an event parameter, such as a user name. The exact string is matched with the given parameter and no regular expression is supported.

Match Versus Strict-Match

The `match` and `strict-match` keywords control the filter behavior for incidents whose event definition does not contain all the parameters specified in a `configure log filter events match`

command. This is best explained with an example. Suppose an event in the XYZ component, named XYZ.event5, contains a physical port number, a source MAC address, but no destination MAC address. If you configure a filter to match a source MAC address and a destination MAC address, XYZ.event5 will match the filter when the source MAC address matches regardless of the destination MAC address, since the event contains no destination MAC address. If you specify the strict-match keyword, then the filter will never match, since XYZ.event5 does not contain the destination MAC address.

In other words, if the match keyword is specified, an incident will pass a filter so long as all parameter values in the incident match those in the match criteria, but all parameter types in the match criteria need not be present in the event definition.

More Information

See the command `show log` for more information about severity levels.

To get a listing of the components present in the system, use the following command:

```
show log components
```

To get a listing of event condition definitions, use the following command:

```
show log events
```

To see the current configuration of a filter, use the following command:

```
show log configuration filter {filter name}
```

Example

By default, all log targets are associated with the built-in filter, DefaultFilter. Therefore, the most straightforward way to send additional messages to a log target is to modify DefaultFilter. In the following example, the command modifies the built-in filter to allow incidents in the *STP* component, and all subcomponents of STP, of severity critical, error, warning, notice and info. For any of these events containing a physical port number as a match parameter, limit the incidents to only those occurring on physical ports 3, 4 and 5 on slot 1, and all ports on slot 2:

```
configure log filter DefaultFilter add events stp severity info match ports 1:3-1:5, 2:*
```

If desired, issue the `unconfigure log DefaultFilter` command to restore the DefaultFilter back to its original configuration.

History

This command was first available in ExtremeXOS 10.1.

New parameter *type* values, including *esrp* and *eaps* were added in ExtremeXOS 11.0 and 11.1.

Support for IPv6 addresses was added in ExtremeXOS 11.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure log messages privilege

```
configure log messages privilege [ admin | user ]
```

Description

This command configures the minimum user account level needed to view logs.

Syntax Description

messages	NVRAM and memory-buffer message targets.
privilege	Configure minimum privilege level needed to view logs.
admin	Only admin (read-write) accounts can view log.
user	User (read-only) accounts can view log also (default).

Default

User.

Usage Guidelines

Use this command to configure the account level needed to view logs.

History

This command was first available in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure log target filter

```
configure log target [console | memory-buffer | | primary-node | |  
  backup-node | nvram | session | syslog [all | ipaddress {udp-port  
  {udp_port}} | ipPort | ipaddress tls-port {tls_port} ] {vr vr_name}  
  {local0...local7}] filter filter-name {severity severity {only}}
```

Description

Associates a filter to a target.

In a stack, this command is applicable only to Master and Backup nodes. This command is not applicable to standby nodes.

Syntax Description

target	Specifies the device to send the log entries.
console	Specifies the console display.
memory-buffer	Specifies the switch memory buffer.
primary-node	Specifies the primary node in a stack.
backup-node	Specifies the backup node in a stack.
nvr	Specifies the switch NVRAM.
session	Specifies the current session (including console display).
syslog	Specifies a syslog remote server.
all	Specifies all of the syslog remote servers.
<i>ipaddress</i>	Specifies the syslog IP address.
<i>ipPort</i>	Specifies the UDP port number for the syslog target.
<i>vr_name</i>	Specifies the virtual router that can reach the server IP address. Note: User-created VRs are supported only on the platforms listed for this feature in the Switch Engine 32.3 Feature License Requirements document.
local0 ... local7	Specifies the local syslog facility.
<i>filter-name</i>	Specifies the filter to associate with the target.
<i>severity</i>	Specifies the minimum severity level to send (if the keyword only is omitted).
only	Specifies that only the specified severity level is to be sent.

Default

If severity is not specified, the severity level for the target is left unchanged. If a virtual router is not specified, VR-Mgmt is used.

Usage Guidelines

This command associates the specified filter and severity with the specified target. A filter limits messages sent to a target.

Although each target can be configured with its own filter, by default, all targets are associated with the built-in filter, DefaultFilter. Each target can also be configured with its own severity level. This provides the ability to associate multiple targets with the same filter, while having a configurable severity level for each target.

A message is sent to a target if the target has been enabled, the message passes the associated filter, the message is at least as severe as the configured severity level, and the message output matches the regular expression specified. By default, the memory buffer and NVRAM targets are enabled. For

other targets, use the command `enable log target`. The following table describes the default characteristics of each type of target.

Table 13: Default target log characteristics

Target	Enabled	Severity Level
console display	no	info
memory buffer	yes	debug-data
NVRAM	yes	warning
session	no	info
syslog	no	debug-data

The built-in filter, DefaultFilter, and a severity level of info are used for each new telnet session. These values may be overridden on a per-session basis using the `configure log target filter` command and specify the target as session. Use the following form of the command for per-session configuration changes:

```
configure log target session filter filtername {severity severity
{only}}
```

Configuration changes to the current session target are in effect only for the duration of the session, and are not saved in FLASH memory. The session option can also be used on the console display, if the changes are desired to be temporary. If changes to the console-display are to be permanent (saved to FLASH memory), use the following form of the command:

```
configure log target console filter filtername {severity severity
{only}}
```

SummitStack Only

The backup-node target is only active on the primary-node, and the primary-node target is active on backup-node and standby-nodes.

Example

The following example sends log messages to the previously syslog host at 10.31.8.25, port 8993, and facility local3, that pass the filter myFilter and are of severity warning and above:

```
configure log target syslog 10.31.8.25:8993 local3 filter myFilter severity warning
```

The following example sends log messages to the current session, that pass the filter myFilter and are of severity warning and above:

```
configure log target session filter myFilter severity warning
```

History

This command was first available in ExtremeXOS 10.1.

The *ipPort* parameter was first available in ExtremeXOS 11.1.

The udp-port parameter was added in ExtremeXOS 21.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure log target format

For console display, session, memory buffer, and NVRAM targets:

```
configure log target [ console | session | memory-buffer | nvram ]
  format [timestamp [seconds | hundredths | none]] [date [ dd-Mmm-yyyy
| yyyy-mm-dd | Mmm-dd | mm-dd-yyyy | mm/dd/yyyy | dd-mm-yyyy | none]]
  {event-name [component | condition | none]} {process-name} {severity}
  {source-line} {host-name}
```

For Syslog targets:

```
configure log target syslog [all | ipaddress {udp-port {udp_port}} |
ipPort | ipaddress tls-port {tls_port}] {vr vr_name} {local} format
[timestamp [ seconds | hundredths | none]] [date [ dd-Mmm-yyyy |
yyyy-mm-dd | Mmm-dd | mm-dd-yyyy | mm/dd/yyyy | dd-mm-yyyy | none]]
{event-name [component | condition | none]} {severity} {priority}
{host-name} {source-line} {tag-id} {tag-name}
```

Description

Configures the formats of the displayed message, on a per-target basis.

In a stack, this command is applicable only to Master and Backup nodes and not applicable to the standby nodes.

Syntax Description

console	Specifies the console display.
session	Specifies the current session (including console display).
memory-buffer	Specifies the switch memory buffer.
nvram	Specifies the switch NVRAM.
syslog	Specifies a syslog target.
all	Specifies all remote syslog servers.
<i>ipaddress</i>	Specifies the syslog IP address.
<i>ipPort</i>	Specifies the UDP port number for the syslog target.
<i>vr_name</i>	Specifies the virtual router that can reach the server IP address. Note: User-created VRs are supported only on the platforms listed for this feature in the Switch Engine 32.3 Feature License Requirements document.

local0 ... local7	Specifies the local syslog facility.
timestamp	Specifies a timestamp formatted to display seconds, hundredths, or none.
date	Specifies a date formatted as specified, or none.
event-name	Specifies how detailed the event description will be. Choose from none, component or condition.
host-name	Specifies whether to include the syslog host name.
priority	Specifies whether to include the priority.
process-name	Specifies whether to include the internal process name.
severity	Specifies whether to include the severity.
source-line	Specifies whether to include the source file name and line number.
tag-id	Specifies whether to include the tag ID.
tag-name	Specifies whether to include the tag name.

Default

The following defaults apply to console display, memory buffer, NVRAM, and session targets:

- timestamp—hundredths
- date—mm-dd-yyyy
- event-name—condition
- process-name—off
- severity—on
- source-line—off
- host-name—off

The following defaults apply to syslog targets (per RFC 3164):

- timestamp—seconds
- date—mmm-dd
- event-name—none
- severity—on
- priority—on
- host-name—off
- source-line—off
- tag-id—off
- tag-name—on

If a virtual router is not specified, VR-Mgmt is used.

Usage Guidelines

This command configures the format of the items that make up log messages. You can choose to include or exclude items and set the format for those items, but you cannot vary the order in which the items are assembled.

When applied to the targets console or session, the format specified is used for the messages sent to the console display or telnet session. Configuration changes to the session target, be it either a telnet or console display target session, are in effect only for the duration of the session, and are not saved in FLASH.

When this command is applied to the target memory-buffer, the format specified is used in subsequent `show log` and `upload log` commands. The format configured for the internal memory buffer can be overridden by specifying a format on the `show log` and `upload log` commands.

When this command is applied to the target syslog, the format specified is used for the messages sent to the specified syslog host.

Timestamps

Timestamps refer to the time an event occurred, and can be output in either seconds as described in RFC 3164 (for example, "13:42:56"), hundredths of a second (for example, "13:42:56.98"), or suppressed altogether. To display timestamps as hh:mm:ss, use the seconds keyword, to display as hh:mm:ss.HH, use the hundredths keyword, or to suppress timestamps altogether, use the none keyword. Timestamps are displayed in hundredths by default.

Date

The date an event occurred can be output as described in RFC 3164. Dates are output in different formats, depending on the keyword chosen. The following lists the date keyword options, and how the date "March 26, 2005" would be output:

- Mmm-dd—Mar 26
- mm-dd-yyyy—03/26/2005
- dd-mm-yyyy—26-03-2005
- yyyy-mm-dd—2005-03-26
- dd-Mmm-yyyy—26-Mar-2005

Dates are suppressed altogether by specifying none. Dates are displayed as mm-dd-yyyy by default.

Event Names

Event names can be output as the component name only by specifying event-name component and as component name with condition mnemonic by specifying event-name condition, or suppressed by specifying event-name none. The default setting is event-name condition to specify the complete name of the events.

Host Name

The configured *SNMP* name of the switch can be output as HOSTNAME described in RFC 3164 by specifying host-name. The default setting is off.

Process Name

For providing detailed information to technical support, the (internal) ExtremeXOS task names of the applications detecting the events can be displayed by specifying process-name. The default setting is off.

Severity

A four-letter abbreviation of the severity of the event can be output by specifying severity on or suppressed by specifying severity off. The default setting is severity on. The abbreviations are: Crit, Erro, Warn, Noti, Info, Summ, Verb, and Data. These correspond to: Critical, Error, Warning, Notice, Informational, Debug-Summary, Debug-Verbose, and Debug-Data.

Source Line

For providing detailed information to technical support, the application source file names and line numbers detecting the events can be displayed by specifying source-line. The default setting is off. You must enable debug mode using the `enable log debug-mode` command to view the source line information. For messages generated prior to enabling debug mode, the source line information is not displayed.

Tag ID

The process-id of the (internal) ExtremeXOS process that generated the event that resulted in the log message can be displayed by specifying tag-id. The default setting is off.

Tag Name

The name of the log component to which the generated event belongs can be displayed by specifying tag-name. The default setting is on. The tag name would be the same as the output of event-name component.

Example

In the following example, the switch generates the identical event from the component *SNTP (Simple Network Time Protocol)*, using three different formats.

Using the default format for the session target, an example log message might appear as:

```
05/29/2005 12:15:25.00 <Warn:SNTP.RslvSrvrFail> The SNTP server parameter value
(TheWrongServer.example.com) can not be resolved.
```

If you set the current session format using the following command:

```
configure log target session format timestamp seconds date mm-dd-yyyy event-name component
```

The same example would appear as:

```
05/29/2005 12:16:36 <Warn:SNTP> The SNTP server parameter value
(TheWrongServer.example.com) can not be resolved.
```

To provide some detailed information to technical support, you set the current session format using the following command:

```
configure log target session format timestamp hundredths date mmm-dd event-name condition
source-line process-name
```

The same example would appear as:

```
May 29 12:17:20.11 SNTP: <Warn:SNTP.RslvSrvrFail> tSntpc: (sntpcLib.c:606) The SNTP
server parameter value (TheWrongServer.example.com) can not be resolved.
```

History

This command was first available in ExtremeXOS 10.1.

The ipPort and host-name parameters were first introduced in ExtremeXOS 11.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure log target match

```
configure log target [console | memory-buffer | nvram | primary-node |
backp-node | session | syslog [all | ipaddress {udp-port {udp_port}}
| ipPort | ipaddress tls-port {tls_port} ] {vr vr_name} {local0 ...
local7}] match [any | match-expression]
```

Description

Associates a match expression to a target.

In a stack, this command is applicable only on a Master and Backup nodes. This command is not applicable for standby nodes.

Syntax Description

console	Specifies the console display.
memory-buffer	Specifies the switch memory buffer.
nvr am	Specifies the switch NVRAM.
primary-node	Specifies the primary node in a stack.
backp-node	Specifies the backup-node in a stack.
session	Specifies the current session (including console display).
syslog	Specifies a syslog target.

all	Specifies all of the remote syslog servers.
<i>ipaddress</i>	Specifies the syslog IP address.
<i>ipPort</i>	Specifies the UDP port number for the syslog target.
<i>vr_name</i>	Specifies the virtual router that can reach the server IP address. Note: User-created VRs are supported only on the platforms listed for this feature in the Switch Engine 32.3 Feature License Requirements document.
local0 ... local7	Specifies the local syslog facility.
any	Specifies that any messages will match. This effectively removes a previously configured match expression.
<i>match-expression</i>	Specifies a regular expression. Only messages that match the regular expression will be sent.

Default

By default, targets do not have a match expression. If a virtual router is not specified, [VR-Mgmt](#) is used.

Usage Guidelines

This command configures the specified target with a match expression. The filter associated with the target is not affected. A message is sent to a target if the target has been enabled, the message passes the associated filter, the message is at least as severe as the configured severity level, and the message output matches the regular expression specified.

See the command [show log](#) for a detailed description of simple regular expressions. By default, targets do not have a match expression.

Specifying any instead of match-expression effectively removes a match expression that had been previously configured, causing any message to be sent that has satisfied all of the other requirements.

To see the configuration of a target, use the following command:

```
show log configuration target {console | memory-buffer | nvram |
primary-node | backup-node | session | syslog {ipaddress {udp-port
{udp_port }}| ipPort | ipaddress tls-port {tls_port}} {vr vr_name}
{[local0...local7]}}
```

To see the current configuration of a filter, use the following command:

```
show log configuration filter {filter name}
```

Example

The following command sends log messages to the current session, that pass the current filter and severity level, and contain the string user5:

```
configure log target session match user5
```

History

This command was first available in ExtremeXOS 10.1.

The *ipPort* parameter was first available in ExtremeXOS 11.1.

The **udp-port** parameter was added in ExtremeXOS 21.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure log target memory-buffer alert percent-full

```
configure log target memory-buffer alert percent-full [ percent | none ]
```

Description

This command configures the log buffer threshold alert.

Syntax Description

percent-full	Generate a log event when the memory buffer percentage fully exceeds the specified percentage threshold.
<i>percent</i>	Percent-full threshold to generate a log event [50-100].
none	No alert.

Default

None.

Usage Guidelines

Use this command to configure the log buffer threshold alert.

History

This command was first available in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure log target severity

```
configure log target [console | memory-buffer | nvram | primary-  
node | backup-node | session | syslog [all | ipaddress {udp-port
```



```
{udp_port}} | ipPort | ipaddress tls-port {tls_port} ] {vr vr_name}
{local0...local7 } ] {severity severity {only}}
```

Description

Sets the severity level of messages sent to the target.

In a stack, this command is applicable only to Master and Backup nodes. You cannot run this command on standby nodes.

Syntax Description

console	Specifies the console display.
memory-buffer	Specifies the switch memory buffer.
nvr am	Specifies the switch NVRAM.
primary-node	Specifies the primary node in a stack.
backup-node	Specifies the backup node in a stack.
session	Specifies the current session (including console display).
syslog	Specifies a syslog target.
all	Specifies all of the remote syslog servers.
<i>ipaddress</i>	Specifies the syslog IP address.
<i>ipPort</i>	Specifies the UDP port number for the syslog target.
<i>vr_name</i>	Specifies the virtual router that can reach the server IP address. Note: User-created VRs are supported only on the platforms listed for this feature in the Switch Engine 32.3 Feature License Requirements document.
local0 ... local7	Specifies the local syslog facility.
<i>severity</i>	Specifies the least severe level to send (if the keyword only is omitted).
only	Specifies that only the specified severity level is to be sent.

Default

By default, targets are sent messages of the following severity level and above:

- console display—info
- memory buffer—debug-data
- NVRAM—warning
- session—info
- syslog—debug-data
- primary node—warning (stack only)
- backup node—warning (stack only)

If a virtual router is not specified, *VR-Mgmt* is used.

Usage Guidelines

This command configures the specified target with a severity level. The filter associated with the target is not affected. A message is sent to a target if the target has been enabled, the message passes the associated filter, the message is at least as severe as the configured severity level, and the message output matches the regular expression specified.

See the command `show log` for a detailed description of severity levels.

To see the current configuration of a target, use the following command:

```
show log configuration target {console | memory-buffer | nvram |
primary-node | backup-node | session | syslog {ipaddress {udp-port
{udp_port }}| ipPort | ipaddress tls-port {tls_port}} {vr vr_name}
{[local0...local7]}}
```

To see the current configuration of a filter, use the following command:

```
show log configuration filter {filter name}
```

Example

The following command sends log messages to the current session, that pass the current filter at a severity level of info or greater, and contain the string user5:

```
configure log target session severity info
```

History

This command was first available in ExtremeXOS 10.1.

The *ipPort* parameter was first available in ExtremeXOS 11.1.

The **udp-port** parameter was added in ExtremeXOS 21.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure log target syslog

```
configure log target syslog [all | ipaddress {udp-port {udp_port}}
| ipPort | ipaddress tls-port {tls_port}] {vr vr_name}
{[local0...local7] from source-ip-address
```

Description

This command specifies the source-ip-address to use when sending log messages to the Syslog server. The Syslog server's IP address along with the ipPort and local facility (a tuple) identify which Syslog server target is to be configured.

Syntax Description

syslog	Specifies a Syslog target.
all	Specifies all of the remote Syslog servers.
<i>ipaddress</i>	Specifies the Syslog server's IP address.
udp-port	Remote Syslog server UDP port. Default 514.
<i>udp_port</i>	UDP port number.
<i>ipPort</i>	Specifies the UDP port number for the Syslog target.
tls_port	Specifies remote Syslog server Transport Layer Security (TLS) for connection type.
<i>tls_port</i>	TLS port number.
<i>vr_name</i>	Specifies the virtual router that can reach the server IP address. Note: User-created VRs are supported only on the platforms listed for this feature in the Switch Engine 32.3 Feature License Requirements document.
local10 ... local17	Specifies the local Syslog facility.
<i>source-ip-address</i>	Specifies the local source IP address to use. Note: The address family (i.e IPv4 or IPv6) of the specified source IP address must be the same as the address family of the Syslog server's.

Default

If a virtual router is not specified, the following virtual routers are used:

- ExtremeXOS 10.1—VR-0
- ExtremeXOS 11.0 and later—[VR-Mgmt](#)

Usage Guidelines

Use this command to identify and configure the Syslog server's IP address. By configuring a source IP address, the Syslog server can identify from which switch it received the log message.

If you do not configure a source IP address for the Syslog target, the switch uses the IP address in the configured VR that has the closed route to the destination.

Example

The following command configures the IP address for the specified Syslog target:

```
configure log target syslog 10.12.1.15 from 10.234.56.78
configure log target syslog 2001:12:1::1 from 2001:44::1
```

History

This command was first available in ExtremeXOS 10.1.

The **udp-port** parameter and support for the EMS to send log messages to Syslog servers having IPv6 address was added in ExtremeXOS 21.1.

Transport Layer Security (TLS) option added in ExtremeXOS 22.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure log target upm filter

```
configure log target upm {upm_profile_name} filter filter-name {severity
[[severity] {only}]}
```

Description

Configures a log target to receive events that conform to a specific EMS filter and severity level requirements.

Syntax Description

<i>upm_profile_name</i>	Specifies a UPM log target to configure.
<i>filter-name</i>	Assigns an EMS filter to the specified log target.
<i>severity</i>	Specifies the minimum severity level for events sent to the log target.
only	Specifies that only events at the specified severity are sent to the log target.

Default

N/A.

Usage Guidelines

Events that meet the criteria established in the EMS filter and the optional severity requirements are forwarded to the UPM log target profile. You can further restrict the forwarded events with the following command:

```
configure log target upm {upm_profile_name} match {any | regex}.
```

Example

The following example configures UPM log target testprofile1 to receive events that meet the criteria defined in EMS filter testfilter1:

```
configure log target upm testprofile1 filter testfilter1
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the Universal Port feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure log target upm match

```
configure log target upm {upm_profile_name} match {any | regex}
```

Description

Configures a log target to receive only those events that meet the specified match criteria.

Syntax Description

<i>upm_profile_name</i>	Specifies the UPM log target to be configured.
any	Matches any event. Use this option to remove a limitation configured with the <i>regex</i> option.
<i>regex</i>	Specifies an expression that must be contained in all forwarded events.

Default

N/A.

Usage Guidelines

This command further restricts the events selected by the command: `configure log target upm {upm_profile_name} filter filter-name {severity [[severity] {only}]}`.

Example

The following example configures UPM log target testprofile1 to receive events that meet the criteria contain the text warning:

```
configure log target upm testprofile1 match warning
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the Universal Port feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure log target xml-notification filter

```
configure log target xml-notification xml_target_name filter filter-name
    {severity [[severity] {only}]}
```

Description

Configures a Web server target with an EMS filter.

Syntax Description

<i>xml_target_name</i>	Specifies the name of the xml notification target.
<i>filter-name</i>	Specifies the name of the EMS filter.
<i>severity</i>	Specifies the least severe level to send (if the keyword only is omitted).

Default

N/A.

Usage Guidelines

Use this command to configure a Web server target with an EMS filter. All EMS filters can be applied.

Example

The following command configures the Web server target test2 with EMS filter filtertest2:

```
configure log target xml-notification test filter filtertest2
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure mac-lockdown-timeout ports aging-time

```
configure mac-lockdown-timeout ports [all | port_list] aging-time  
seconds
```

Description

Configures the MAC address lock down timeout value in seconds for the specified port or group of ports or for all ports on the switch.

Syntax Description

all	Specifies all ports.
<i>port_list</i>	Specifies one or more ports or slots and ports.
<i>seconds</i>	Configures the length of the time out value in seconds. The default is 15 seconds; the range is 15 to 2,000,000 seconds.

Default

The default is 15 seconds.

Usage Guidelines

This timer overrides the *FDB* aging time.

This command only sets the duration of the MAC address lock down timer. To enable the lock down timeout feature, use the following command:

```
enable mac-lockdown-timeout ports [all | port_list]
```

Example

The following command configures the MAC address lock down timer duration for 300 seconds for ports 2:3, 2:4, and 2:6:

```
configure mac-lockdown-timeout ports 2:3, 2:4, 2:6 aging-time 300
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure mac-locking ports first-arrival aging

```
configure mac-locking ports port_list first-arrival aging [enable |  
disable]
```

Description

Enables and disables the aging of first-arrival MAC addresses.

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
enable	MAC addresses aged out from the forwarding database are removed from MAC locking.
disable	MAC addresses aged out from the forwarding database are not removed from MAC locking.

Default

First-arrival MAC lock aging is disabled by default.

Usage Guidelines

This command does not apply to MAC addresses locked by static locking.

When enabled, first-arrival MAC addresses that are aged out of the forwarding database are removed from the associated port MAC lock. New MAC addresses can be learned until the configured first-arrival limit is reached.

Example

The following command enables first-arrival MAC lock aging on port 2:3:

```
configure mac-locking ports 2:3 first-arrival aging enable
```

History

This command was first available in ExtremeXOS 15.7.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure mac-locking ports first-arrival limit-learning

```
configure mac-locking ports port_list first-arrival limit-learning  
learn_limit
```


Description

Configures dynamic MAC locking on a port by restricting MAC locking on a port to a maximum number of end station addresses first connected to that port.

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
<i>learn_limit</i>	Specifies the maximum number of first-arrival end station MAC addresses that can be connected to the port. Valid values are 0–600.

Default

600 first-arrival end station MAC addresses

Usage Guidelines

When the configured limit is reached, no further entries are learned. If, however, the learned entries are aged out, new MAC addresses can be learned.

You cannot specify a value that is lower than the number of MACs locked in the MAC lock station table.

Example

The following example configures 400 as the maximum number of first-arrival MAC addresses that can connect to port 14.

```
configure mac-locking ports 14 first-arrival limit-learning 400
```

History

This command was first available in ExtremeXOS 15.7.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure mac-locking ports first-arrival link-down-action

```
configure mac-locking ports port_list first-arrival link-down-action  
[clear-macs | retain-macs]
```

Description

Clears or retains first arrival MAC locking addresses when the link goes down.

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
clear-macs	First arrival MAC locking addresses will be cleared when the link goes down.
retain-macs	First arrival MAC locking addresses will be retained when the link goes down.

Default

When the link goes down, by default, all the first arrival MAC locking addresses will be removed (cleared).

Usage Guidelines

If you specify **retain-macs**, the first arrival MAC locking addresses will be retained even when the link goes down.

Example

The following example disables the clearing of first arrival MAC locking addresses on port 14.

```
configure mac-locking ports 14 first-arrival link-down-action retain-macs
```

History

This command was first available in ExtremeXOS 15.7.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure mac-locking ports first-arrival move-to-static

```
configure mac-locking ports port_list first-arrival move-to-static
```

Description

Moves all current first-arrival MAC locking addresses to static entries.

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
------------------	---

Default

N/A

Usage Guidelines

This command converts dynamic MAC locked station entries to static MAC locked entries. The static MAC locked entries are saved in configuration and preserved across reboots.

This command does not convert the forwarding database entries to static-permanent entries.

Example

The following example converts the dynamic MAC locked station entries on port 14 to static MAC locked entries.

```
configure mac-locking ports 14 first-arrival move-to-static
```

History

This command was first available in ExtremeXOS 15.7.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure mac-locking ports learn-limit-action

```
configure mac-locking ports port_list learn-limit-action [disable-port |  
remain-enabled]
```

Description

Configures a port to be disabled or remain enabled when the port learns the configured maximum number of MACs.

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
disable-port	Disables the port when the configured MAC limit is reached.
remain-enabled	Port remains enabled after the configured MAC limit is reached.

Default

The port remains enabled after the configured MAC limit is reached.

Usage Guidelines

This command is used for both first arrival and static MAC locking methods.

Example

The following example configures port 14 to be disabled when the configured MAC limit is reached.

```
configure mac-locking ports 14 learn-limit-action disable-port
```

History

This command was first available in ExtremeXOS 15.7.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure mac-locking ports log

```
configure mac-locking ports port_list log {violation | threshold} [on | off]
```

Description

Enables or disables the sending of a syslog message for MAC lock messages.

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
violation	Sends a syslog message if the maximum value configured for dynamic and static MAC locking is exceeded.
threshold	Sends a syslog message if the maximum value configured for dynamic and static MAC locking is reached.
on	Sending a syslog message for the specified event is enabled.
off	Sending a syslog message for the specified event is disabled.

Default

If neither **violation** nor **threshold** is specified, violation is used by default.

Usage Guidelines

When MAC locking violations are enabled, the device sends a syslog message if a connected end station exceeds the maximum value configured for dynamic and static MAC locking.

When MAC locking thresholds are enabled, the device sends an syslog message if a connected end station reaches the maximum value configured for dynamic and static MAC locking.

Example

The following example enables threshold syslog messages on port 14.

```
configure mac-locking ports 14 log threshold on
```

History

This command was first available in ExtremeXOS 15.7.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure mac-locking ports static delete station

```
configure mac-locking ports port_list static delete station  
[station_mac_address | all]
```

Description

Deletes MAC locking for all MAC address or the specified MAC address on the specified port.

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
<i>station_mac_address</i>	Specifies the MAC address from which MAC locking will be deleted.
all	Deletes MAC locking from all MAC addresses associated with the specified port.

Default

N/A

Usage Guidelines

None.

Example

The following example deletes MAC locking from the MAC address 00-a0-c9-0d-32-11 on port 14.

```
configure mac-locking ports 14 static delete station 00-a0-c9-0d-32-11
```

History

This command was first available in ExtremeXOS 15.7.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure mac-locking ports *port_list* **static limit-learning** *learn_limit*

```
configure mac-locking ports port_list static limit-learning learn_limit
```

Description

Restricts MAC locking on a port to a maximum number of static (management defined) MAC addresses for end stations connected to this port.

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
<i>learn_limit</i>	Specifies the maximum number of static end station MAC addresses that can be connected to the port. Valid values are 0-64.

Default

64 static end station MAC addresses.

Usage Guidelines

When the configured limit is reached, no further entries are learned. If, however, the learned entries are aged out, new MAC addresses can be learned.

You cannot set a value that is lower than the number of MACs locked in the MAC lock station table.

You cannot configure the learning limit on both a port and a port-VLAN. If the learning limit is configured on a port, configuration on a port-VLAN will not be allowed. Similarly, if the learning limit is configured on a port-VLAN, configuration on port is not allowed.

Example

The following example configures 40 as the maximum number of static MAC addresses that can connect to port 14.

```
configure mac-locking ports 14 static limit-learning 40
```

History

This command was first available in ExtremeXOS 15.7.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure mac-locking ports static

```
configure mac-locking ports port_list static [add | enable | disable]  
                        station station_mac_address
```

Description

Creates, enables, and disables a static MAC locking entry.

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
add	Adds a MAC locking association between the specified MAC address and port.
enable	Enables an existing MAC locking association between the specified MAC address and port.
disable	Disables an existing MAC locking association between the specified MAC address and port.
<i>station_mac_address</i>	Specifies the MAC address.

Default

A static MAC locking association is enabled by default.

Usage Guidelines

Up to 64 MAC addresses can be locked per port.

When added and enabled, a static MAC lock configuration allows only the end station designated by the MAC address to participate in frame relay.

Disabled entries are counted when calculating the total number of locked stations.

Example

The following example creates a MAC locking association between port 14 and 00-a0-c9-0d-32-11.

```
configure mac-locking ports 14 static add 00-a0-c9-0d-32-11
```

History

This command was first available in ExtremeXOS 15.7.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure mac-locking ports trap

```
configure mac-locking ports port_list trap {violation | threshold} [on | off]
```

Description

Enables or disables the sending of an *SNMP* trap for MAC lock messages.

Syntax Description

violation	Sends an SNMP trap if the maximum value configured for dynamic and static MAC locking is exceeded.
threshold	Sends an SNMP trap if the maximum value configured for dynamic and static MAC locking is reached.
on	Sending an SNMP trap for the specified event is enabled.
off	Sending an SNMP trap for the specified event is disabled.

Default

If neither **violation** nor **threshold** is specified, violation is used by default.

Usage Guidelines

When MAC locking violations are enabled, the device sends an SNMP trap if a connected end station exceeds the maximum value configured for dynamic and static MAC locking.

When MAC locking thresholds are enabled, the device sends an SNMP trap if a connected end station reaches the maximum value configured for dynamic and static MAC locking.

Example

The following example enables threshold traps on port 14.

```
configure mac-locking ports 14 trap threshold on
```

History

This command was first available in ExtremeXOS 15.7.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure macsec cipher-suite

```
configure macsec cipher-suite [gcm-aes-128 | gcm-aes-256] ports
    port_list
```

Description

Configures the preferred cipher suite for MAC Security (MACsec).

Syntax Description

cipher-suite	Selects provisioning MACsec cipher suite to be used if elected as key server.
gcm-aes-128	Galois/Counter Mode of AES-128 symmetric block cipher (Default).
gcm-aes-256	Galois/Counter Mode of AES-256 symmetric block.
ports	Specifies configuring ports.
<i>port_list</i>	Lists which ports to configure the selected cipher suite on.

Default

The cipher suite **gcm-aes-128** is selected by default.

Usage Guidelines

Table 14: Cipher Support

GCM-AES-256 and GCM-AES-128	
Ports with LRM/MACsec Adapter	
ExtremeSwitching 5320, 5420, 5720 on all ports. ExtremeSwitching 5520 on all ports, except 5520-VIM-4X and 24X 10G ports.	

If GCM-AES-256 is desired between two switches using the LRM/MACsec Adapter, you need to issue this command on at least the key server side, but preferably on both sides.

If the port is elected as MKA key server, then the configured cipher suite is used to protect all port traffic. If the peer port is elected as MKA key server, then the peer chooses which cipher suite to use.

Example

The following example selects the gcm-aes-256 cipher suite on ports 22, 30-33:

```
# configure macsec cipher-suite gcm-aes-256 22,30-33
```

The following example selects the gcm-aes-128 cipher suite on port 30:

```
# configure macsec cipher-suite gcm-aes-128 30
```

History

This command was first available in ExtremeXOS 30.2.

Platform Availability

This command is available on the following platforms.



Note

The MACsec feature requires the installation of the MAC Security feature pack license.

Platform	Ports
ExtremeSwitching 5320	All ports of all models except stacking ports.
ExtremeSwitching 5420	All ports of all models except stacking ports.
ExtremeSwitching 5520	All ports, except 5520-VIM-4X and 5520-24X 10G ports
ExtremeSwitching 5720	All ports of all models except stacking ports.

configure macsec connectivity-association

```
configure macsec connectivity-association ca_name [pre-shared-key {ckn
ckn} {cak [encrypted encrypted_cak] | cak} | ports [port_list]
[enable | disable]]
```

Description

Configures a previously created connectivity-association (CA) object that holds MAC Security (MACsec) key authentication data. For a particular CA, you can change the pre-shared key and enable/disable authentication on one or more ports.

Syntax Description

connectivity-association	Secures connectivity provided between MACsec stations.
<i>ca_name</i>	Selects CA object to configure.
pre-shared-key	Selects static MACsec key consisting of both a CKN and CAK:
ckn	Selects changing the CA key name. This public (non-secret) key name allows each of the MKA participants to select which connectivity association key (CAK) to use to process a received MACsec key agreement (MKA) protocol packets (MKPDU).
<i>ckn</i>	Sets the CA key name. Length allowed is 1-32 characters, entered as ASCII or an octet string preceded with 0x.

cak	Sets the connectivity association key (CAK). If you are using 256-bit cipher suite, then the CAK must be 32 octets. The 128-bit cipher suite can use either a 16- or 32-octet CAK. This is a long-lived secret key used to derive short-lived lower-layer keys (ICK, KEK, and SAK) that are used for key distribution and data encryption.
<i>cak</i>	Sets the non-encrypted CAK value. Must be entered as an octet string (for example: "0x859e72f0..."). A 128-bit (16 octet) CAK requires 32 hexadecimal digits, and a 256-bit (32 octet) CAK requires 64 hexadecimal digits. These values are secret and should be generated off switch with a suitable pseudorandom number generator.
encrypted	Designates that secret key value is in encrypted format.
<i>encrypted_cak</i>	Sets the value for the secret key. The encrypted CAK value is generated by the <code>show configuration macsec</code> command for previously configured CAKs.
ports	Specifies configuring ports.
<i>port_list</i>	Lists which ports to configure.
enable	Enable the MKA connectivity association on the selected port list.
disable	Disables the MKA connectivity association on the selected port list.

Default

N/A.

Usage Guidelines

You can only enable/disable CAs on ports that support MACsec.

If execution of this command results in MACsec being enabled on more than 48 ports for a given 5320 or 5420 series switch, then the command will fail.

Example

The following example sets CKN to "the red key" and CAK to a 128-bit key "0x01020304050607080910111213141516" for CA object "testca":



Note

The CAK shown here is an example. Use your own random number for maximum security.

```
configure macsec connectivity-association testca pre-shared-key ckn "the red key" cak
"0x01020304050607080910111213141516"
```

The following example enables MACsec authentication on port 13 for CA object "testca":

```
# configure macsec connectivity-association testca ports 13 enable
```

The following example disables MACsec authentication on port 13 for CA object "testca":

```
# configure macsec connectivity-association testca ports 13 disable
```

History

This command was first available in ExtremeXOS 30.1.

Support for 256-cipher suite was added in ExtremeXOS 30.2.

Platform Availability

This command is available on the following platforms.



Note

The MACsec feature requires the installation of the MAC Security feature pack license.

Platform	Ports
ExtremeSwitching 5320	All ports of all models except stacking ports.
ExtremeSwitching 5420	All ports of all models except stacking ports.
ExtremeSwitching 5520	All ports, except 5520-VIM-4X and 5520-24X 10G ports
ExtremeSwitching 5720	All ports of all models except stacking ports.

configure macsec include-sci

```
configure macsec include-sci [enable | disable] ports port_list
```

Description

Configures the include-SCI flag to ensure interoperability with third-party devices that do not decode encrypted MAC Security (MACsec) packets when the SCI is not present.

Syntax Description

include-sci	Provision inclusion of SCI in SecTAG field while transmitting MACsec frames.
enable	Include SCI in SecTAG.
disable	Do not include SCI in SecTAG (Default).
ports	Specifies configuring ports.
<i>port_list</i>	Lists which ports to configure the include-SCI flag on.

Default

Disabled by default (SCI is not included in MAC Security Tag (SecTAG)).

Usage Guidelines

The SecTAG appended to each data packet contains an optional parameter called Secure Channel Indicator (SCI). The SCI is used to identify the sending Secure Association (SA) when the connectivity-association (CA) comprises three or more peers.

Because ExtremeXOS only supports point-to-point links (which have exactly two peers), the SCI is not sent by default (which saves 8-octets per SecTAG'd packet). Certain third-party MACsec devices, such as the CentOS's MACsec client and Cisco Catalyst 3650, fail to decode encrypted MACsec packets when the SCI is not present. To ensure interoperability with such devices, you can configure the Include-SCI flag. When this flag is set, the port always includes the 8-octet SCI in the SecTAG of all outgoing packets.



Important

After enabling MACsec, if you change the include-SCI flag, you must run the `configure macsec initialize ports port_list` command afterward. Otherwise, the change is not applied.

Example

The following example enables including SCI in SecTAG field while transmitting MACsec frames on port 13:

```
configure macsec include-sci enable port 13
```

The following example disables including SCI in SecTAG field while transmitting MACsec frames on port 44:

```
# configure macsec include-sci disable port 44
```

History

This command was first available in ExtremeXOS 30.1.

Platform Availability

This command is available on the following platforms.



Note

The MACsec feature requires the installation of the MAC Security feature pack license.

Platform	Ports
ExtremeSwitching 5320	All ports of all models except stacking ports.
ExtremeSwitching 5420	All ports of all models except stacking ports.
ExtremeSwitching 5520	All ports, except 5520-VIM-4X and 5520-24X 10G ports
ExtremeSwitching 5720	All ports of all models except stacking ports.

configure macsec initialize ports

```
configure macsec initialize ports port_list
```

Description

Resets the MAC Security (MACsec) Key Agreement (MKA) protocol state machine on one or more ports and applies MACsec configuration changes to already enabled ports.

Syntax Description

initialize	Selects resetting the MACsec Key Agreement protocol state machine.
ports	Specifies configuring ports.
<i>port_list</i>	Lists which ports to reset the MACsec Key Agreement protocol state machine on.

Default

N/A.

Usage Guidelines

Issuing this command resets the MKA state machine, which in turn deletes any secured channels and their secure association keys (SAKs). This command is also used to apply MACsec configuration changes (mka actor-priority, include-sci, replay-protect, mka life-time) to an already enabled port. All traffic is blocked until MKA renegotiates a new set of keys and those keys are installed. For more information, see *IEEE802.1X-2010 Clause 12.9.3 Initialization*.

Example

The following example resets the MACsec Key Agreement protocol state machine on port 13:

```
configure macsec initialize ports 13
```

History

This command was first available in ExtremeXOS 30.1.

Platform Availability

This command is available on the following platforms.



Note

The MACsec feature requires the installation of the MAC Security feature pack license.

Platform	Ports
ExtremeSwitching 5320	All ports of all models except stacking ports.
ExtremeSwitching 5420	All ports of all models except stacking ports.
ExtremeSwitching 5520	All ports, except 5520-VIM-4X and 5520-24X 10G ports
ExtremeSwitching 5720	All ports of all models except stacking ports.

configure macsec mka actor-priority

```
configure macsec mka actor-priority actor_priority ports port_list
```

Description

Configures MAC Security (MACsec) actor's priority for port(s).

Syntax Description

mka	Configures MACsec key agreement (MKA) parameters.
actor-priority	Designates setting the priority advertised during MKA key server election.
<i>actor-priority</i>	Sets the actor priority value. A lower value denotes higher priority. Range is 0-255 or 0x0-0xFF. Default is 0x10.
ports	Specifies configuring ports.
<i>port_list</i>	Lists which ports to configure the actor priority on.

Default

Default value for actor priority is 0x10.

Usage Guidelines

Each MKA participant selects the participant advertising the highest priority as the key server. In the event of a tie, the participant with the highest priority MAC address (lowest value) is selected. The recommended priority range for infrastructure ports is 0x00 to 0x1f, with a default of 0x10. You can assign the full range of priorities, 0x00 to 0xff:

- To *have* a port become a key server, raise the priority by assigning a priority value *less* than 0x10.

- To *not* have a port become key server, lower the priority by assigning a priority value *greater* than 0x10.



Important

After enabling MACsec, if you change the actor priority, you must run the `configure macsec initialize ports port_list` command afterward. Otherwise, the change is not applied.

Example

The following example raises the actor priority value to 0x5 on port 13:

```
# configure macsec mka actor-priority "0x5" port 13
# configure macsec initialize port 13
```

The following example lowers the actor priority value to "31" on port 14:

```
# configure macsec mka actor-priority 31 port 14
# configure macsec initialize port 14
```

History

This command was first available in ExtremeXOS 30.1.

Platform Availability

This command is available on the following platforms.



Note

The MACsec feature requires the installation of the MAC Security feature pack license.

Platform	Ports
ExtremeSwitching 5320	All ports of all models except stacking ports.
ExtremeSwitching 5420	All ports of all models except stacking ports.
ExtremeSwitching 5520	All ports, except 5520-VIM-4X and 5520-24X 10G ports
ExtremeSwitching 5720	All ports of all models except stacking ports.

configure macsec mka life-time

```
configure macsec mka life-time mka_life_time ports port_list
```

Description

Configures MAC Security (MACsec) lifetime for port(s).

Syntax Description

mka	Configures MACsec key agreement (MKA) parameters.
life-time	Designates setting the lifetime of potential and live peers. Expiration causes removal from a list, and higher intervals increase MKA protocol stability.
<i>mka_life_time</i>	Sets the lifetime of potential and live peers. Range is 6-30. Default is 6 seconds.
ports	Specifies configuring ports.
<i>port_list</i>	Lists which ports to configure the actor priority on.

Default

Default value for life-time 6 seconds.

Usage Guidelines

If MACsec link flap occurs, loosen the `life-time` equally on both sides of the MACsec connection.



Note

MACsec link flap is likely to only occur on links connected to lower-end switches (the ExtremeSwitching X620 switch, for example).



Important

After enabling MACsec, if you change the MKA lifetime, you must run the `configure macsec initialize ports port_list` command afterward. Otherwise, the change is not applied.

Example

The following configures the MKA lifetime to 10 seconds on port 3:

```
# configure macsec mka life-time 10 port 3
# configure macsec initialize port 3
```

History

This command was first available in ExtremeXOS 31.5.

Platform Availability

This command is available on the following platforms:



Note

The MACsec feature requires the installation of the MAC Security feature pack license.

Platform	Ports
ExtremeSwitching 5320	All ports of all models except stacking ports.
ExtremeSwitching 5420	All ports of all models except stacking ports.
ExtremeSwitching 5520	All ports, except 5520-VIM-4X and 5520-24X 10G ports
ExtremeSwitching 5720	All ports of all models except stacking ports.

configure macsec replay-protect

```
configure macsec replay-protect [window_size_in_packets | disable] ports
port_list
```

Description

Configures MAC Security (MACsec) replay-protect window size for port(s).

Syntax Description

replay-protect	Configures dropping out-of-order packets received on a port.
<i>window_size_in_packets</i>	Sets replay-protect window size value. Out-of-order packets up to selected value are accepted. Range is 0–4,294,967,295. Default is 0 (out-of-order packets are dropped).
disable	Disables replay protection. Out-of-order packets are allowed.
ports	Specifies configuring ports.
<i>port_list</i>	Lists which ports to configure the replay-protect window on.

Default

Default value for replay-protect window is 0 packets, which drops all out-of-order packets.

Usage Guidelines

The replay protection feature provides for the dropping of out-of-order packets received on a port. The window size is set to 0 by default, meaning any packet received out-of-order is dropped. Setting the window size to non-zero sets the range of sequence numbers that are tolerated, to allow receipt of

packets that have been misordered by the network. If replay protection is disabled, packet sequence numbers are not checked and out-of-order packets are not dropped.



Important

After enabling MACsec, if you change the replay protect window size, you must run the `configure macsec initialize ports port_list` command afterward. Otherwise, the change is not applied.

Example

The following example disables replay protection on port 13:

```
# configure macsec replay-protect disable port 13
# configure macsec initialize port 13
```

The following example sets replay-protect window size to 50 packets on port 14. If the last data packet received has a packet number (PN) of N, then the next received packet is accepted if its PN is greater than or equal to N-50. If the PN is less than N-50, the packet is dropped and the "Late Pkts" counter is incremented:

```
# configure macsec replay-protect 50 port 14
# configure macsec initialize port 14
```

History

This command was first available in ExtremeXOS 30.1.

Platform Availability

This command is available on the following platforms.



Note

The MACsec feature requires the installation of the MAC Security feature pack license.

Platform	Ports
ExtremeSwitching 5320	All ports of all models except stacking ports.
ExtremeSwitching 5420	All ports of all models except stacking ports.
ExtremeSwitching 5520	All ports, except 5520-VIM-4X and 5520-24X 10G ports
ExtremeSwitching 5720	All ports of all models except stacking ports.

configure mcast ipv4 cache timeout

```
configure mcast ipv4 cache timeout {seconds | none}
```

Description

Configures the IPv4 multicast cache timeout.

Syntax Description

<i>seconds</i>	Idle time after which cache entries are deleted.
none	Cache entries are not timed out.

Default

300 seconds.

Usage Guidelines

Cache timeout is the time after which the cache entries are deleted if traffic is not received for that duration. This applies only for snooping and MVR caches and does not apply for PIM caches.

The range is 90 to 100000 seconds. You can use the option **none** if you do not want the cache entry to be deleted. If **none** is configured, the cache entries can be deleted only using the following command:

```
clear igmp snooping
```

Example

The following example configures the IPv4 multicast cache timeout to 400 seconds.

```
configure mcast ipv4 cache timeout 400
```

The following command clears the IPv4 multicast cache timeout.

```
configure mcast ipv4 cache timeout none
```

History

This command was first available in ExtremeXOS 15.2.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the *IGMP* snooping feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure mcast ipv6 cache timeout

```
configure mcast ipv6 cache timeout {seconds | none}
```

Description

Configures the IPv6 multicast cache timeout.

Syntax Description

<i>seconds</i>	Idle time after which cache entries are deleted.
none	Cache entries are not timed out.

Default

300 seconds.

Usage Guidelines

Cache timeout is the time after which the cache entries are deleted if traffic is not received for that duration. This applies only for snooping and MVR caches and does not apply for PIM caches.

The range is 90 to 100000 seconds. You can use the option **none** if you do not want the cache entry to be deleted. If **none** is configured, the cache entries could be deleted only using the following command:

```
clear igmp snooping
```

Example

The following example configures the IPv6 multicast cache timeout to 400 seconds.

```
configure mcast ipv6 cache timeout 400
```

The following command clears the IPv6 multicast cache timeout.

```
configure mcast ipv6 cache timeout none
```

History

This command was first available in ExtremeXOS 15.2.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MLD snooping feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure meter

```
configure meter metername {committed-rate cir [Gbps | Mbps | Kbps
| Pps]} {max-burst-size burst-size [Kb | Mb | packets]} {out-
actions [{disable-port} {drop | set-drop-precedence {dscp [none |
dscp-value]} {dot1p [none | dot1p-value ]}}] {log} {trap}}] {ports
[port_group | port_list]}
```

Description

Configures an ACL meter to provide ingress traffic rate shaping.

Syntax Description

<i>metername</i>	Specifies the ACL meter name.
committed-rate	Specifies the committed information rate in gigabits per second (Gbps), megabits per second (Mbps), or kilobits per second (Kbps).
max-burst-size	Specifies the maximum burst size or peak burst size in kilobits (Kb) or megabits (Mb).
out-actions	Specifies actions to take if traffic exceeds the profile.
drop	Specifies to drop out of profile traffic.
set-drop-precedence	Specifies to mark packet for high drop precedence.
dscp	Specifies to set DSCP.
<i>dscp-value</i>	DSCP value (0-63).
none	Specifies to leave the DSCP or dot1p value unchanged.
dot1p	Specifies dot1p value to be set.
<i>dot1p-value</i>	Dot1p value (0-7).
log	Generate log event if traffic exceeds configured rate.
trap	Generate <i>SNMP</i> trap if traffic exceeds configured rate.
ports	Meter configuration is applicable to ports in the specified port_group or port_list.
<i>port_group</i>	Port group name.
<i>port_list</i>	Port list separated by a comma.

Default

By default, a newly committed meter has no maximum burst size, no committed rate, and a default action of drop.

Usage Guidelines

The meter configured with this command is associated with an ACL rule by specifying the meter name using the meter action modifier within the rule.

The committed-rate keyword specifies the traffic rate allowed for this meter, and the configured rate operates as described in [Table 15](#). The rate you specify is rounded up to the next granularity increment value. For example, if you configure a 1 Mbps committed rate for a platform with a 64Kbps granularity increment, this value falls between the increment values of 960 Kbps and 1024 Kbps, so the effective

committed rate is set to 1024 Kbps. Also, note that some platforms listed below require an adjustment to the expected rate to calculate the configured rate.

Table 15: Rate Configuration Notes

Platform	Granularity	Notes
All platforms	64Kbps	Specify the traffic rate in Kbps, Mbps, or Gbps. The range is 64Kbps to 1 Gbps for GE ports and 1Mbps to 10 Gbps for 10GE ports. Add 20 bytes per frame to the expected rate to determine the configured rate.

The `max-burst-size` keyword specifies the maximum number of consecutive bits that are allowed to be in-profile at wire-speed. The `max-burst-size` parameter can be specified in Kb, Mb, or Gb. The specified `max-burst-size` is rounded down to the nearest supported size. The `max-burst-size` range on ExtremeSwitching switches is 32Kb to 128Mb.

The keyword `out-actions` specifies the action that is taken when a packet is out-of-profile. The supported actions include dropping the packet, marking the drop precedence for the packet, setting the DSCP value in the packet, or setting the DOT1P value in the packet. The keyword `drop` indicates that any out-of-profile packet is immediately dropped. The keyword `set-drop-precedence` marks out-of-profile packets with high drop precedence. If the optional keyword `set-dscp` is specified, the DSCP value, as specified by the parameter `dscp-value`, is written into the out-of-profile packet. Setting the DSCP value to `none` leaves the DSCP value in the packet unchanged. If the optional keyword `set-dot1p` is specified, the DOT1P value, as specified by the parameter `dot1p-value`, is written into the out-of-profile packet. Setting the DOT1P value to `none` leaves the DOT1P value in the packet unchanged.

Example

The following example configures the ACL meter `maximum_bandwidth`, assigns it a rate of 10 Mbps, and sets the out of profile action to drop:

```
configure meter maximum_bandwidth committed-rate 10 Mbps out-action drop
```

The following example uses the `port_groups` variable:

```
configure meter ingmeter0 committed-rate 50 Mbps out-actions drop log disable-port ports
GroupA
configure meter ingmeter1 committed-rate 75 Mbps out-actions drop log disable-port ports
GroupA
configure meter ingmeter0 committed-rate 100 Pps out-actions drop log disable-port ports
GroupB
configure meter ingmeter1 committed-rate 150 Pps out-actions drop log disable-port ports
GroupB
```

History

This command was available in ExtremeXOS 11.1.

The **log**, **trap** and **ports** keywords and `port-group` and `port_list` variables were added in ExtremeXOS 16.1

The **dot1p** keyword and variable were added in ExtremeXOS 21.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure mirror add

```
configure mirror { mirror_name } add [ {vlan} vlan_name | vlan vlan_id]
    {ingress | [port port {ingress}] | ip-fix | port port vlan [vlan_id |
    vlan_name ] {ingress}]
```

Description

Specifies mirror source filters for an instance.

Syntax Description

<i>mirror_name</i>	Specifies the mirror's name.
vlan	Specifies a <u>VLAN</u> .
<i>vlan_id</i>	Specifies a VLAN ID.
port	Specifies a port or slot and port.
<i>port</i>	Specifies particular ports or slots and ports.
ingress	Specifies packets be mirrored as they are received on a port. Note: This parameter is available only with port-based mirroring.
ip-fix	Enables mirroring of the first fifteen packets of every IPFIX flow.
egress	Specifies packets be mirrored as they are sent from a port. Note: This parameter is available only with port-based mirroring.
ingress-and-egress	Specifies all forwarded packets be mirrored. This is the default for port-based mirroring. Note: This parameter is available only with port-based mirroring.

Default

N/A.

Usage Guidelines

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

Port mirroring configures the switch to copy all traffic associated with one or more ports to a monitor port on the switch. The switch uses a traffic filter that copies a group of traffic to the monitor port.

Up to 128 mirroring filters can be configured with the restriction that a maximum of 16 of these can be configured as VLAN and/or virtual port (port + VLAN) filters.

One monitor port or 1 monitor port list can be configured. A monitor port list may contain up to 16 ports.

Frames that contain errors are not mirrored.

For general guideline information and information for various platforms, see “Guidelines for Mirroring” in the [Switch Engine 32.3 User Guide](#) or the Usage Guidelines of the [enable mirroring to port](#) command.

Example

The following example sends all traffic coming into a switch on port 11 and the VLAN default to the mirror port:

```
configure mirror add port 11 vlan default
```

History

This command was first available in ExtremeXOS 15.3.

The `vlan_id` option was added in ExtremeXOS 16.1.

The **ip-fix** option was added in ExtremeXOS 21.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure mirror add ports anomaly

```
configure mirror add ports port_list anomaly
```

Description

Mirrors detected anomaly traffic to the mirror port.

Syntax Description

<i>port_list</i>	Specifies the list of ports.
------------------	------------------------------

Default

N/A.

Usage Guidelines

The command mirrors detected anomaly traffic to the mirror port. You must enable a mirror port and enable protocol anomaly protection on the slot that has the port to be monitored before using this command. After configuration, only detected anomaly traffic from these ports are dropped or mirrored to the mirror port, and legitimate traffic is not affected.

This command takes effect after enabling anomaly-protection.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure mirror control_index

```
configure mirror control_index [ add | delete ] mirror_name
```

Description

Adds or deletes existing mirrors to a mirror MIB instance (specified by a control index) .

Syntax Description

<i>mirror_name</i>	Specifies a specific mirror name to add to or delete from a mirror MIB instance.
<i>control_index</i>	Mirror destination control index (1-4). Also known as: etsysMirrorDestinationControlIndex. Each comprises a group of mirror names.
add	Specifies adding a mirror name to group referenced by a control index.
delete	Specifies deleting a mirror name from a group referenced by a control index.

Default

N/A.

Usage Guidelines

To use policy-based mirroring, you need a mirror MIB instance (designated by a control index) with one or more associated mirrors to apply mirrors to a policy profile.

Only mirrors with a single 'to' port or remote-ip can be applied to a mirror MIB instance.

Example

The following example adds existing mirror "mirror1" to mirror MIB instance with control index "2":

```
configure mirror 2 add mirror1
```

History

This command was first available in ExtremeXOS 30.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure mirror delete

```
configure mirror {mirror_name} delete [ {vlan} vlan_name | vlan vlan_id]
      {port port} | ip-fix | port port vlan [vlan_id | vlan_name]
```

Description

Deletes mirror source filters for an instance.

Syntax Description

<i>mirror_name</i>	Specifies the mirror's name.
<i>port</i>	Specifies a port or a slot and port.
<i>port</i>	Specifies particular ports or slots and ports.
ip-fix	Disables mirroring packets of IPFIX flows.
vlan	Specifies a <u>VLAN</u> .
<i>vlan_id</i>	Specifies a VLAN ID.
<i>name</i>	Specifies a VLAN name.

Default

N/A.

Example

The following example deletes the mirroring filter on port 1:

```
configure mirroring delete ports 1
```

History

This command was first available in ExtremeXOS 10.1.

The VLAN option was added in ExtremeXOS 11.0.

The `vlan_id` option was added in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure mirror description

```
configure mirror mirror_name description [ mirror-desc | none ]
```

Description

Creates, edits or deletes a mirroring instance description string.

Syntax Description

<i>mirror_name</i>	Specifies the mirror name.
description	Specifies the mirror description to create or edit.
none	Deletes the existing mirror description.

Default

N/A.

Usage Guidelines

Use this command to create, edit or delete a mirroring instance description string.

Example

The following example configures the mirror description.

```
configure mirror description
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure mirror name

```
configure mirror mirror_name name new_name
```

Description

Updates or specifies the "to port" definitions for a named mirroring instance .

Syntax Description

<i>mirror_name</i>	Specifies the mirror name.
name	Specifies a new mirror name.
<i>new_name</i>	Specifies the new mirror name.

Default

N/A.

Usage Guidelines

Use this command to update or specify the "to port" definitions for a named mirroring instance.

Example

```
configure mirror m1 name m2
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure mirror to

```
configure mirror mirror_name {to [port port | port-list port_list  
| loopback port port] | remote-ip {add} remote_ip_address [{vr}  
vr_name] } {from [source_ip_address | auto-source-ip] } {ping-check [on  
| off]}} {remote-tag rtag | port none} {priority priority_value}
```

Description

Updates or specifies the "to port", "to port list", or remote IP address destination definitions for a named mirroring instance.

Syntax Description

<i>mirror_name</i>	Specifies the mirror name.
port	Specifies the mirror output port.
port-list	Specifies the list of ports where traffic is to be mirrored.
loopback port	Specifies an otherwise unused port required when mirroring to a port_list. The loopback-port is not available for switching user data traffic.
<i>port</i>	Specifies a single loopback port that is used internally to provide this feature.
remote-tag	Specifies the value of the <i>VLAN</i> ID used by the mirrored packets when egressing the monitor port.
port	Specifies the port definition for the mirroring instance.
none	Specifies none for the to port definition.
remote-ip	Sends mirrored packets to specified remote destination IP address.
<i>remote_ip_address</i>	Specifies the destination remote IP address for mirrored packets.
add	Adds a redundant (more than one) remote IP address with a unique priority to a mirror instance.
vr	Specifies the virtual router of the remote IP address.
<i>vr_name</i>	Specifies the virtual router name. If not specified, VR of current command context is used.
from	Configures source IP address of encapsulated mirrored packets.
<i>source_ip_address</i>	Specifies the local source IPv4 address for encapsulated mirrored packets.
auto-source-ip	Automatically use source IP address of egress VLAN to be used to reach remote IP address.
ping-check	Configure ping health check for remote IP address.
on	Only send mirrored packets to remote IP address if periodic pings to remote IP address are successful (default).
off	Send mirrored packets to remote IP address without any ping health check, assuming MAC address and port of next hop IP address are static or learned.
priority	Configures a unique priority value for each redundant remote IP address of a mirror instance.
<i>priority_value</i>	Sets a unique priority value for a remote IP address. The priority value must be unique for each remote IP address in the mirror instance. The range is from 1 (least preferred) to 100 (most preferred). The default is 50.

Default

Ping health check of the remote IP address is enabled unless otherwise specified.

If a VR is not specified, the VR of the current command context is used.

The default priority value for multiple redundant IP addresses is 50.

Usage Guidelines

Use this command to update, or specify the "to port", "to port-list", or remote IP address destination definitions for a named mirroring instance.

The **none** keyword can be used to remove a previously configured port/port-list, or remote IP address on a disabled mirror instance.

For high availability, you can add up to four redundant remote IP addresses. For each mirror instance, the remote IP address with the highest configured priority value that has status "up" is used as the destination IP address for GRE-tunneled mirrored traffic. All other remote IP addresses deemed "up" for that mirror instance are standby—ready to be used in the event the preferred remote IP address becomes "down". If you are adding another (redundant) remote IP address to an existing mirror that already has a remote IP address configured, you must use the **add** option.

The **remote-ip** cannot be the IP from vr-mgmt.

You cannot specify vr-mgmt as VR.

Example

The following example configures a mirror instance to port 3, slot 4:

```
# configure mirror to port 3:4
```

The following example configures multiple (redundant) remote IP addresses ("5.1.1.2", "4.1.1.2", "3.1.1.2", "2.1.1.2") for mirror "analytics_chicago_1":

```
# enable mirror analytics_chicago_1 to remote-ip 5.1.1.2
# configure mirror analytics_chicago_1 to remote-ip add 4.1.1.2 priority 40
# configure mirror analytics_chicago_1 to remote-ip add 3.1.1.2 priority 30
# configure mirror analytics_chicago_1 to remote-ip add 2.1.1.2 priority 20
# configure mirror analytics_chicago_1 add vlan v1

# show mirror

analytics_chicago_1    (Enabled)
  Description:
    Mirror to remote IP: 5.1.1.2          VR          : VR-Default
    From IP              : Auto source IP  Ping check: On
    Priority              : 50
    Status                : Up. Active

    Mirror to remote IP: 4.1.1.2          VR          : VR-Default
    From IP              : Auto source IP  Ping check: On
    Priority              : 40
    Status                : Up. Standby

    Mirror to remote IP: 3.1.1.2          VR          : VR-Default
    From IP              : Auto source IP  Ping check: On
    Priority              : 30
    Status                : Down. Ping timed out

    Mirror to remote IP: 2.1.1.2          VR          : VR-Default
    From IP              : Auto source IP  Ping check: On
```

```

Priority          : 20
Status           : Up. Standby
Source filter instances used : 1
    All ports, vlan vl, ingress only

```

History

This command was first available in ExtremeXOS 15.3.

The remote IP address option was added in ExtremeXOS 22.4.

Redundant remote IP addresses capability was added in ExtremeXOS 30.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure mirror to remote-ip delete

```

configure mirror {mirror_name to remote-ip delete [all |
    remote_ip_address [{vr} vr_name]} }

```

Description

Removes one or all of the redundant remote IP addresses from a mirror instance.

Syntax Description

<i>mirror_name</i>	Mirror instance name.
to	Selects mirroring to another location.
remote-ip	Send mirrored packets to specified destination IP address using L2 GRE encapsulation.
delete	Delete all or one remote IP addresses from a mirror instance.
all	Delete all remote IP addresses from a disabled mirror instance.
<i>remote_ip_address</i>	Delete the specified existing remote IP address from a mirror instance.
vr	Specifies a virtual router.
<i>vr_name</i>	Specifies the name of the virtual router to which this command applies. If a name is not specified, the current CLI context is used.

Default

If a virtual router is not specified, the current CLI context is used.

Usage Guidelines

To delete all or the last remaining remote IP address, you must disable the mirror first (`disable mirror [mirror_name | all]`).

Example

The following example removes the remote IP address "1.1.3.3" from the mirror instance "m1":

```
# configure mirror m1 to remote-ip delete 1.1.3.3
```

History

This command was first available in ExtremeXOS 30.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure mirror to remote-ip protocol-type

```
configure mirror to remote-ip protocol-type [erspan-v1 | trans-ether-bridging | user-defined protocol_value]
```

Description

Adds a configurable GRE protocol type for mirror-to-remote IP addresses.

Syntax Description

mirror	Specifies configuring mirrors.
to	Selects mirroring to another location.
remote-ip	Sends mirrored packets to specified destination IP address using L2 GRE encapsulation.
protocol-type	Selects GRE protocol type in the header of mirrored packets to all remote IP addresses.
erspan-v1	Specifies GRE protocol type 0x88BE, Encapsulated Remote Switched Port Analyzer version 1, also known as ERSPAN type II (default).
trans-ether-bridging	Specifies GRE protocol type 0x6558, Trans Ether Bridging.
user-defined	Specifies GRE protocol type specified in hexadecimal (for example, 0x6558).
<i>protocol_value</i>	Specifies a user-defined, two-byte hexadecimal value for GRE protocol type (for example, 0x6558).

Default

By default, the type is **erspan-v1**.

Usage Guidelines

The configured value is global, and the new value is applied immediately in hardware for all active mirrors to remote IP addresses.

To view the current setting, use the `show mirror [mirror_name | control_index | mirror_name_li] | [all | enabled]` command.

Example

The following example sets the type as **trans-ether-bridging**:

```
# configure mirror to remote-ip protocol-type trans-ether-bridging
```

History

This command was first available in ExtremeXOS 30.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure mlag peer alternate ipaddress

```
configure mlag peer peer_name alternate ipaddress ip_address vr vr_name
| none
```

Description

This command configures the IP address for alternate health check mechanism.

Syntax Description

<code>mlag</code>	Multi-switch Link Aggregation used to combine remote ports and local ports to a common logical connection.
<code>peer</code>	Multi-switch Link Aggregation Group peer switch.
<code>peer_name</code>	Alphanumeric string identifying the peer.
<code>alternate</code>	Health check on an alternate path.
<code>ipaddress</code>	MLAG peer IP address for alternate path health checks.
<code>vr</code>	Virtual router.
<code>vr_name</code>	Virtual router name.
<code>none</code>	Do not use alternate path health checks.

Default

None.

Usage Guidelines

Use this command to configure the IP address for alternate health check mechanism. Use the **none** option to unconfigure the configured IP.

Example

The following example displays `show mlag peer` output with the alternate path IP configured:

```
w4.10 # show mlag peer
Multi-switch Link Aggregation Peers:

MLAG Peer      : sw3
VLAN           : two
Local IP Address : 2100:51:2::4
Peer IP Address  : 2100:51:2::3
Virtual Router  : VR-Default
MLAG ports     : 1
Tx-Interval    : 100 ms
Checkpoint Status : Up
Peer Tx-Interval : 100 ms
Rx-Hellos      : 13212
Tx-Hellos      : 13485
Rx-Checkpoint Msgs: 121
Tx-Checkpoint Msgs: 316
Rx-Hello Errors : 0
Tx-Hello Errors : 0
Hello Timeouts  : 0
Checkpoint Errors : 0
Up Time        : 0d:0h:17m:47s
Peer Conn.Failures: 0
Local MAC      : 00:04:96:51:ac:d7
Peer MAC       : 00:04:96:36:52:91
Config'd LACP MAC : None
Current LACP MAC : 00:04:96:51:ac:d7
Authentication:  : md5
Authentication Key: .{:OFarc#'qX)+6zid#smIE+',+)ocijk (encrypted)

Alternate path information:
VLAN           : Mgmt
Local IP Address : 10.127.7.74
Peer IP Address  : 10.127.7.73
Virtual Router  : VR-Mgmt
Rx-Hellos      : 243
Tx-Hellos      : 551
Rx-Hello Errors : 0
Tx-Hello Errors : 0
Hello Timeouts  : 1
```

When the alternate path IP is not configured, the following output is shown:

```
sw4.10 # show mlag peer
Multi-switch Link Aggregation Peers:

MLAG Peer      : sw3
VLAN           : two
Local IP Address : 2100:51:2::4
Peer IP Address  : 2100:51:2::3
Virtual Router  : VR-Default
MLAG ports     : 1
Tx-Interval    : 100 ms
Checkpoint Status : Up
Peer Tx-Interval : 100 ms
Rx-Hellos      : 13212
Tx-Hellos      : 13485
Rx-Checkpoint Msgs: 121
Tx-Checkpoint Msgs: 316
Rx-Hello Errors : 0
Tx-Hello Errors : 0
Hello Timeouts  : 0
Checkpoint Errors : 0
Up Time        : 0d:0h:17m:47s
Peer Conn.Failures: 0
Local MAC      : 00:04:96:51:ac:d7
Peer MAC       : 00:04:96:36:52:91
Config'd LACP MAC : None
Current LACP MAC : 00:04:96:51:ac:d7
Authentication:  : md5
Authentication Key: .{:OFarc#'qX)+6zid#smIE+',+)ocijk (encrypted)

Alternate path information: None
```

History

This command was first available in ExtremeXOS 15.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure mlag peer authentication

```
configure mlag peer peer_name authentication [md5 key {encrypted
encrypted_auth_key | auth_key } | none]
```

Description

Configures the MD5 authentication key for checkpoint connection to peer.

Syntax Description

mlag	Multi-switch Link Aggregation Group used to combine remote ports and local ports to a common logical connection.
peer	Multi-switch Link Aggregation Group peer switch.
<i>peer_name</i>	Alphanumeric string identifying the MLAG peer.
authentication	Authentication for MLAG checkpoint connection.
md5	MD5 authentication type.
key	Authentication key for checkpoint connection to the MLAG peer.
encrypted	Authentication key is in encrypted format.
<i>auth_key</i>	Authentication key. Max 32 characters.
none	Do not use authentication.

Default

None.

Usage Guidelines

Use this command to configure MD5 authentication key for checkpoint connection to MLAG peer.

Example

The following example displays `show mlag peer` output when authentication is not configured:

```
* Switch # show mlag peer
Multi-switch Link Aggregation Peers:

MLAG Peer      : p2
VLAN           : isc
Virtual Router : VR-Default
```

```

Local IP Address : 10.1.1.1      Peer IP Address : 10.1.1.2
MLAG ports      : 1             Tx-Interval    : 1000 ms
Checkpoint Status : Up          Peer Tx-Interval : 1000 ms
Rx-Hellos       : 8722          Tx-Hellos      :
8725

Rx-Checkpoint Msgs: 1322        Tx-Checkpoint Msgs: 947
Rx-Hello Errors  : 0           Tx-Hello Errors  : 0
Hello Timeouts   : 0           Checkpoint Errors : 0
Up Time          : 0d:2h:22m:26s Peer Conn.Failures: 0
Local MAC        : 00:04:96:7e:13:93 Peer MAC         : 00:04:96:7e:13:71
Config'd LACP MAC : None        Current LACP MAC  : 00:04:96:7e:13:71
Authentication   : None

Alternate path information:
VLAN              : Mgmt          Virtual Router    : VR-Mgmt
Local IP Address  : 1.1.1.1      Peer IP Address   : 1.1.1.2
Rx-Hellos        : 722          Tx-Hellos        : 725
Rx-Hello Errors   : 0           Tx-Hello Errors   : 0
Hello Timeouts    : 0

```

The following example displays show mlag peer output when authentication is configured:

```

* Switch # show mlag peer
Multi-switch Link Aggregation Peers:

MLAG Peer      : p2
VLAN           : isc
Local IP Address : 10.1.1.1      Virtual Router    : VR-Default
MLAG ports     : 1             Peer IP Address   : 10.1.1.2
Checkpoint Status : Up          Tx-Interval      : 1000 ms
Rx-Hellos      : 8722          Peer Tx-Interval  : 1000 ms
8725            Tx-Hellos      :

Rx-Checkpoint Msgs: 1322        Tx-Checkpoint Msgs: 947
Rx-Hello Errors  : 0           Tx-Hello Errors  : 0
Hello Timeouts   : 0           Checkpoint Errors : 0
Up Time          : 0d:2h:22m:26s Peer Conn.Failures: 0
Local MAC        : 00:04:96:7e:13:93 Peer MAC         : 00:04:96:7e:13:71
Config'd LACP MAC : None        Current LACP MAC  : 00:04:96:7e:13:71
Authentication   : md5
Authentication Key: abcdefghijklmnopqrstuvwxyz (encrypted)

Alternate path information:
VLAN              : Mgmt          Virtual Router    : VR-Mgmt
Local IP Address  : 1.1.1.1      Peer IP Address   : 1.1.1.2
Rx-Hellos        : 722          Tx-Hellos        : 725
Rx-Hello Errors   : 0           Tx-Hello Errors   : 0
Hello Timeouts    : 0

```

History

This command was first available in ExtremeXOS 15.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure mlag peer interval

```
configure mlag peer peer_name interval msec
```

Description

Configures the length of time between health check hello packets.

Syntax Description

<i>peer_name</i>	Specifies an alpha numeric string identifying the peer.
<i>msec</i>	Specifies an MLAG peer health-check hello interval in milliseconds. The range is 50-10000ms. The default is 1000ms.

Default

The interval default is 1000 milliseconds.

Usage Guidelines

Use this command to configure the length of time between health check hello packets exchanged between MLAG peer switches. After three health check hellos are lost, the MLAG peer switch is declared to be failed, triggering an MLAG topology change.

Example

The following command sets an interval of 700 milliseconds on the switch101 peer. switch:

```
# configure mlag peer switch101 interval 700
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure mlag peer ipaddress

```
configure mlag peer peer_name ipaddress peer_ip_address {vr VR}
```

Description

Associates an peer switch with an MLAG peer structure.

Syntax Description

<i>peer_name</i>	Specifies an alpha numeric string identifying the MLAG peer.
<i>peer_ip_address</i>	Specifies an IPv4 or IPv6 address.
VR	Specifies a virtual router.

Default

N/A.

Usage Guidelines

Use this command to associate an MLAG peer structure with an MLAG peer switch IP address.

The specified IP address must be contained within an existing direct route. If not, the following error message is displayed:

ERROR: Specified IP address is not on directly attached subnet in VR.

The link connecting MLAG peer switches should use load sharing. If it does not, a output similar to the following is displayed:

Note: VLAN v1 will be used as the Inter-Switch Connection to the MLAG peer mp1. Warning: The VLAN v1 does not have a load share port configured yet. It is recommended that the Inter-Switch Connection use load sharing.

Example

The following command associates the MLAG peer structure switch101 with the MLAG peer switch IP address 1.1.1.1 on VR-USER:

```
# configure mlag peer switch101 ipaddress 1.1.1.1 vr "VR-USER"
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure mlag peer lacp-mac

```
configure mlag peer peer_name lacp-mac [auto | lacp_mac_address]
```

Description

Configures LACP MAC on each of the MLAG peer switches. This MAC address will be used as the system identifier in the LACPDUs sent over the MLAG ports.

Syntax Description

mlag	Multi-switch link aggregation used to combine remote ports and local ports to a common logical connection.
<i>peer_name</i>	Alphanumeric string identifying the MLAG peer.
lACP-mac	MAC address to be used as the system identifier in LACPDU for MLAG ports.
auto	System identifier in LACPDU automatically uses switch MAC of MLAG peer with higher IP address for ISC control VLAN (default).
<i>lACP_mac_address</i>	MAC address.

Default

Auto.

Usage Guidelines

This command is used to configure the System Identifier used in LACPDU for MLAG ports. The same value has to be configured on both the MLAG peers.

Example

```
# configure mlag peer "peer1" lACP-mac auto
# configure mlag peer "peer1" lACP-mac 00:01:02:03:04:05
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure mlag peer name

```
configure { mlag peer } peer_name name new_peer_name
```

Description

Renames an established peer.

Syntax Description

mlag	Specifies configuring MLAG settings.
peer	Specifies configuring aspects of the MLAG peer switch.
<i>peer_name</i>	Current MLAG peer name.

name	Specifies renaming the MLAG peer.
<i>new_peer_name</i>	Specifies the new name for the MLAG peer.

Default

N/A.

Usage Guidelines

To view changes made with this command, use the `show mlag peer {peer_name}` command.

Example

The following example changes the MLAG peer name from "mlag1" to "mlag2":

```
# configure mlag peer mlag1 name mlag2
```

```
# show mlag peer
Multi-switch Link Aggregation Peers:

MLAG Peer      : mlag2
VLAN           :
Local IP Address :
MLAG ports     : 0
Checkpoint Status : Down
Rx-Hellos      :
Rx-Checkpoint Msgs:
Rx-Hello Errors :
Hello Timeouts : N/A
Up Time        : N/A
Local MAC      : 00:04:96:9b:f5:cc
Config'd LACP MAC : None
Authentication : None

Virtual Router :
Peer IP Address :
Tx-Interval    : 1000 ms
Peer Tx-Interval : N/A ms
Tx-Hellos      :
Tx-Checkpoint Msgs:
Tx-Hello Errors :
Checkpoint Errors :
Peer Conn.Failures: N/A
Peer MAC       : None
Current LACP MAC : 00:04:96:9b:f5:cc

Alternate path information: None
```

History

This command was first available in ExtremeXOS 32.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure mlag ports convergence-control

```
configure mlag ports convergence-control [conserve-access-lists | fast]
```

Description

Sets a preference for having a fast convergence time or conserving access lists.

Syntax Description

conserve-access-lists	Specifies that conserving access lists is preferred over low traffic convergence time.
fast	Specifies that low traffic convergence time is preferred at the expense of the number of user access lists.

Default

Conserve-access-lists.

Usage Guidelines

Achieving fast convergence times on local port state changes (down and up), independent of the number of *FDB* entries learned on the MLAG port, requires the use of ACLs. This limits the number of ACLs you have available. This command allows you to set your preference for having either fast convergence time or conserving available access lists for your users.



Note

Configuring fast convergence-control limits the number of ACLs that can be supported by the switch. You must ensure that the system has sufficient user ACLs free when fast mode is selected. Configuring conserve-access-lists convergence-control may increase convergence times on MLAG port failures.

Fast convergence configuration has global significance in that it applies to all MLAG groups that are currently configured and those that may be configured in the future.

Example

The following command specifies a priority of conserving access lists over low traffic convergence time:

```
# configure mlag ports convergence-control conserve-access-lists
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure mlag ports link-up-isolation

```
configure mlag ports link-up-isolation [on | off]
```

Description

Configures linkup isolation, which prevents flood traffic received on newly operational MLAG ports from being forwarded to ISC ports before the ISC blocking filter is installed.

Syntax Description

on	Isolate MLAG ports from sending traffic to local ISC port during link-up transition until remote ISC port is configured.
off	Do not isolate MLAG ports from sending traffic to local ISC port during link-up transition.

Default

The default is off.

Usage Guidelines

Under certain circumstances, a temporary (less than a second) loop condition exists when an MLAG port becomes operational, but before the remote MLAG peer installs the ISC blocking filter. MLAG linkup isolation addresses this condition by preventing any flood traffic (broadcast, unknown, unicast, etc.) received on a just operational MLAG port from being forwarded to ISC ports until the remote MLAG peer installs the ISC blocking filter.

Example

The following example enables MLAG linkup isolation:

```
configure mlag ports link-up-isolation on
```

History

This command was first available in ExtremeXOS 22.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure mlag ports reload-delay

```
configure mlag ports reload-delay reload-delay
```

Description

This command configures a reload delay on Multi-switch Link Aggregation Group (MLAG) ports.

Syntax Description

reload-delay	Specifies creating a reload delay on MLAG ports.
<i>reload-delay</i>	Specifies the MLAG port reload-delay timer in seconds (range = 1-1,200 seconds). The default is 30 seconds.

Default

The default reload-delay timer interval is 30 seconds.

Usage Guidelines

There are cases where MLAG ports comes up quicker than ISC ports after a switch reboot causing traffic loss during this time gap. This command allows you to configure a time delay for MLAG ports providing enough time for ISC ports/neighborship of other Layer 3 protocols to come up. To have this delay timer take effect, you need to issue the [enable mlag port reload-delay](#) on page 2215 command.

To stagger the bringing up of MLAG ports, use the command `configure mlag ports reload-interval [none | reload_interval_msec]`

To view the current selection for reload delay, use the `show mlag ports {port_list}` command.

Example

The following example sets the reload-delay to 60 seconds:

```
# configure mlag ports reload-delay 60
```

History

This command was first available in ExtremeXOS 22.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure mlag ports reload-interval

```
configure mlag ports reload-interval [none | reload_interval_msec]
```

Description

Configures a staggered bringing up of ports.

Syntax Description

reload-interval	Specifies configuring the time between bringing up individual MLAG ports when reload delay is enabled.
none	Specifies not waiting between bringing up individual MLAG ports (default).
<i>reload_interval_msec</i>	Specifies the time interval between bringing up MLAG ports in milliseconds. The range is 0-10,000.

Default

By default, this feature is disabled.

Usage Guidelines

MLAG reload delay timer is used to disable MLAG ports during configuration load to allow time for the convergence of protocols and for reachability of MLAG peers (`configure mlag ports reload-delay reload-delay`). When there is a large number of MLAG ports (50+), and when all of them are brought up at the same time after the reload delay timer expires, a high convergence time of 1.5 seconds might occur. This command configures a time delay between each of the MLAG ports coming up.

To view the current selection for reload interval, use the `show mlag ports {port_list}` command.

Example

The following example configures reload delay interval of 50 milliseconds:

```
# configure mlag ports reload-interval 50
```

History

This command was first available in ExtremeXOS 22.7.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure mld

```
configure mld query_interval query_response_interval
               last_member_query_interval {{vlan} vlan_name} {{vr} vr_name}
               {robustness}
```

Description

Configures the Multicast Listener Discovery (MLD) timers.

Syntax Description

<i>query_interval</i>	Specifies the interval (in seconds) between general queries.
<i>query_response_interval</i>	Specifies the maximum query response time (in seconds).
<i>last_member_query_interval</i>	Specifies the maximum group-specific query response time (in seconds).

<i>vlan_name</i>	Applies the configuration only to the specified <i>VLAN</i> . If no VLAN is specified, the configuration applies to all VLANs.
<i>vr_name</i>	Specifies the VR to which the configuration should be applied. If not parameter is specified, the configuration is applied to the current VR context.
robustness	Specifies the degree of robustness for the network.

Default

- query interval—125 seconds
- query response interval—10 seconds
- last member query interval—1 second
- robustness—2

Usage Guidelines

Timers are based on RFC2710. 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.
- robustness—The degree of robustness of the network. The range is 2 to 7.

Example

The following command configures the MLD timers:

```
configure mld 100 5 1 3
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 multicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure mld snooping fast-learning

```
configure mld snooping fast-learning [on | off] [vlan vlan_name]
```

Description

Configures fast-learning mode.

Syntax Description

<code>vlan_name</code>	Specifies a vlan name
------------------------	-----------------------

Default

off.

Usage Guidelines

When MLD snooping is enabled on a [VLAN](#), learning of group entries will happen only when the next periodic query is sent by the querier in the network. When fast-learning is turned on using this command, a query is sent under the following conditions:

- When MLD snooping is enabled.
- When MLD snooping VLAN is operationally up.
- Group join limit changed through configuration.

Query generated for faster learning uses unspecified address as the source address (both L2 and L3), unless the switch generating the triggered query is the querier for the network.

Example

```
configure mld snooping fast-learning on
```

History

This command was first available in ExtremeXOS 15.2.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MLD snooping feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure mld snooping filters

```
configure mld snooping filters [per-port | per-vlan]
```

Description

Selects the type of MLD snooping filters that are installed.

Syntax Description

per-port	Installs the per-port MLD snooping filters.
per-vlan	Installs the per- <u>VLAN</u> MLD snooping filters.

Default

per-port.

Usage Guidelines

Use the per-vlan option when the number of VLANs configured on the switch is lower than half of the maximum numbers listed in [Table 9](#) on page 584. This option conserves usage of the hardware Layer 3 multicast forwarding table.

When the number of configured VLANs is larger than half of the maximum values listed in [Table 9](#) on page 584, select the per-port option. Each VLAN requires additional interface hardware ACL resources. The per-port option conserves usage of the interface hardware ACL resources.

To display the MLD snooping filters configuration, use the show mld snooping command.

Example

The following command configures the switch to install the per-VLAN MLD snooping filters:

```
configure mld snooping filters per-vlan
```

History

This command was first available in ExtremeXOS 15.2.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MLD snooping feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure mld snooping flood-list

```
configure mld snooping flood-list [policy | none]
```

Description

Configures certain multicast addresses to be slow path flooded within the VLAN.

Syntax Description

<i>policy</i>	Specifies a policy file with a list of multicast addresses to be handled.
none	Specifies no policy file is to be used.

Default

None.

Usage Guidelines

With this command, you can configure certain multicast addresses to be slow path flooded within the VLAN, instead of fast path forwarded according to MLD and/or Layer 3 multicast protocol.

A policy file is a text file with the extension .pol. It can be created or edited with any text editor. The specified policy file *policy file* should contain a list of addresses that determine if certain multicast streams are to be treated specially. Typically, if the switch receives a stream with a destination address which is in the *policy file* in 'permit' mode, that stream is software flooded and no hardware entry is installed.

When adding an IPv6 address into the policy file, a 128-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 a certain stream as control packets.

To create a policy file for the snooping flood-list, use the following template:

```
# This is a template for MLD Snooping Flood-list Policy File
# Add your group addresses between "Start" and "End"
# Do not touch rest of file!!!!
entry mldFlood {
  if match any {
    #----- Start of group addresses -----
    nlri ff05::100:1/128;
    nlri ff05::100:15/128;
    #----- end of group addresses -----
  } then {
    permit;
  }
}
entry catch_all {
  if {
  } then {
    deny;
  }
}
```



Note

The switch does not validate any IP address in the policy file used in this command. Therefore, slow-path flooding should be used only for streams that 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 MLD or PIM), so it should be used with caution.

Slow-path flooding occurs within the L2 VLAN only.

Use the **none** option to effectively disable slow path flooding.

You can use the `show mld` command to see the configuration of slow path flooding.



Note

This command has no effect in the current release, as IPv6 multicast traffic floods on all platforms.

Example

The following example configures the multicast data stream specified in `access1` for slow-path flooding:

```
configure mld snooping flood-list access1
```

The following command specifies that no policy file is to be used, thus effectively disabling slow-path flooding:

```
configure mld snooping flood-list none
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 multicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure mld snooping leave-timeout

```
configure mld snooping leave-timeout leave_timeout_ms [{vlan} vlan_name]  
                        [{vr} vr_name]
```

Description

Configures the MLD snooping leave timeout.

Syntax Description

<i>leave_timeout_ms</i>	Specifies an MLD leave timeout value in milliseconds upon receiving an MLD done message.
vlan_name	Applies the configuration only to the specified <i>VLAN</i> . If no VLAN is specified, the configuration applies to all VLANs.
vr_name	Specifies the VR to which the configuration should be applied. If no parameter is specified, the configuration is applied to the current VR context.

Default

1000 ms.

Usage Guidelines

The range is 0-175000 ms (175 seconds). For timeout values of one second or less, you must set the leave-timeout to a multiple of 100 ms. For values of more than one second, you must set the leave-timeout to a multiple of 1000 ms (one second).

The specified time is the maximum leave timeout value. The switch could leave sooner if an MLD done message is received before the timeout occurs.

Example

The following example configures the MLD snooping leave timeout to 10 seconds:

```
configure mld snooping leave-timeout 10000
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 multicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure mld snooping timer

```
configure mld snooping timer router_timeout host_timeout {{vlan}  
  vlan_name} {{vr} vr_name}
```

Description

Configures the MLD snooping timers.

Syntax Description

<i>router_timeout</i>	Specifies the time in seconds before removing a router snooping entry.
<i>host_timeout</i>	Specifies the time in seconds before removing a host's group snooping entry.
vlan_name	Applies the configuration only to the specified <i>VLAN</i> . If no VLAN is specified, the configuration applies to all VLANs.
vr_name	Specifies the VR to which the configuration should be applied. If no parameter is specified, the configuration is applied to the current VR context.

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 maximum time, in seconds, that a router snooping entry can stay without receiving a router report. The range is 10 to 214,748,364 seconds (6.8 years). The default setting is 260 seconds.
- *host_timeout*—The maximum time, in seconds, that a group snooping entry can stay without receiving a group report. The range is 10 to 214,748,364 seconds (6.8 years). The default setting is 260 seconds.

MLD snooping is a Layer 2 function of the switch. It does not require multicast routing to be enabled. The feature reduces the flooding of IPv6 multicast traffic. On the VLAN, MLD 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 (FF02::x).

MLD snooping is enabled by default on the switch. MLD snooping expects at least one device on every VLAN to periodically generate MLD query messages. Without an MLD querier, the switch eventually stops forwarding IPv6 multicast packets to any port, because the MLD snooping entries times out, based on the value specified in host timeout.

Example

The following example configures the MLD snooping timers to 600 seconds for both timers:

```
configure mld snooping timer 600 600
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 multicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure mld snooping vlan ports add dynamic group

```
configure mld snooping {vlan} vlan_name {ports portlist} add dynamic  
group [IPv6_grp_ipaddress]
```

Description

Configures an MLD dynamic group.

Syntax Description

<i>vlan_name</i>	Specifies a VLAN name.
<i>portlist</i>	Specifies a port list.
<i>IPv6_grp_ipaddress</i>	Specifies the multicast group IPv6 address.

Default

N/A.

Usage Guidelines

This command adds MLD groups to specific VLANs or to ports belonging to specific VLANs. After the groups are added, the expiration timer is started; this causes the groups to expire. The configuration is not saved in the configuration file. The following message is displayed on execution of this command:

```
INFO: This command is not saved in the configuration.
```

Example

The following example configures a dynamic MLD entry so the multicast group ff02::1:1 is forwarded to VLAN marketing on ports 2:1-2:4:

```
configure mld snooping marketing ports 2:1-2:4 add dynamic group ff02::1:1
```

History

This command was first available in ExtremeXOS 15.3.2.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IPv4 multicast feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure mld snooping vlan ports add static group

```
configure mld snooping {vlan} vlan_name {ports port_list } add static  
group IPv6_grp_ipaddress
```

Description

Configures [VLAN](#) ports to receive the traffic from a multicast group, even if no MLD joins have been received on the port.

Syntax Description

<i>vlan_name</i>	Specifies a VLAN name.
<i>port_list</i>	Specifies one or more ports or slots and ports. On a SummitStack, it can be a list of slots (nodes) and ports. On a standalone switch, it can be one or more port numbers. In the form 1, 2, 3-5, 2:5, 2:6-2:8.
<i>IPv6_grp_ipaddress</i>	Specifies the multicast group IPv6 address.

Default

N/A.

Usage Guidelines

Use this command to forward a particular multicast group to VLAN ports. In effect, this command emulates a host on the port that has joined the multicast group. As long as the port is configured with the static entry, multicast traffic for that multicast group is forwarded to that port.

The switch sends proxy MLD messages in place of those generated by a real host. The proxy messages use the VLAN IPv6 address for source address of the messages. If the VLAN has no IPv6 address assigned, the proxy MLD message uses 0::0 as the source IP address.

Example

The following example configures a static MLD entry so the multicast group ff02::1:1 is forwarded to VLAN marketing on ports 2:1-2:4:

```
configure mld snooping marketing ports 2:1-2:4 add static group ff02::1:1
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 multicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure mld snooping vlan ports add static router

```
configure mld snooping {vlan} vlan_name ports port_list add static  
router
```

Description

Configures *VLAN* ports to forward the traffic from all multicast groups, even if no MLD joins have been received on the port.

Syntax Description

<i>vlan_name</i>	Specifies a VLAN name.
<i>port_list</i>	Specifies one or more ports or slots and ports. On a SummitStack, it can be a list of slots and ports. On a standalone switch, it 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 this command to forward all multicast groups to the specified VLAN ports. In effect, this command emulates a multicast router attached to those ports. As long as the ports are configured with the static entry, all available multicast traffic is forwarded to those ports.

Example

The following example configures a static MLD entry so all multicast groups are forwarded to VLAN marketing on ports 2:1-2:4:

```
configure mld snooping marketing ports 2:1-2:4 add static router
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 multicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure mld snooping vlan ports delete static group

```
configure mld snooping {vlan} vlan_name ports port_list delete static  
group [all | v6grpipaddress]
```

Description

Removes the configuration that causes VLAN ports to receive the traffic from a multicast group, even if no MLD joins have been received on the port.

Syntax Description

<i>vlan_name</i>	Specifies a VLAN name.
<i>port_list</i>	Specifies one or more ports or slots and ports. On a modular switch, it can be a list of slots and ports. On a standalone switch, it can be one or more port numbers. In the form 1, 2, 3-5, 2:5, 2:6-2:8.
all	Specifies all multicast groups.
<i>v6grpipaddress</i>	Specifies the multicast group IPv6 address.

Default

N/A.

Usage Guidelines

Use this command to delete a static group from a particular VLAN port.

To add a static group, use the following command:

```
configure mld snooping {vlan} vlan_name portsport_list add static  
groupv6grpipaddress
```

Example

The following example removes a static MLD entry so the multicast group ff02::a:b is not forwarded to VLAN marketing on ports 2:1-2:4, unless an MLD join message is received on the port:

```
configure mld snooping marketing ports 2:1-2:4 delete static group ff02::a:b
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 multicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure mld snooping vlan ports delete static router

```
configure mld snooping {vlan} vlan_name ports port_list delete static  
router
```

Description

Configures VLAN ports to stop forwarding the traffic from all multicast groups, unless MLD joins have been received on the port.

Syntax Description

<i>vlan_name</i>	Specifies a VLAN name.
<i>port_list</i>	Specifies one or more ports or slots and ports. On a SummitStack, it can be a list of slots and ports. On a standalone switch, it can be one or more port numbers. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.

Default

None.

Usage Guidelines

Use this command to remove the configuration that forwards all multicast groups to the specified VLAN ports.

Example

The following example removes a static MLD entry so all multicast groups are not forwarded to VLAN marketing on ports 2:1-2:4, unless an MLD join is received on the port:

```
configure mld snooping marketing ports 2:1-2:4 delete static router
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 multicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure mld snooping vlan ports filter

```
configure mld snooping vlan vlan_name ports port_list filter [policy]
```

Description

Configures a MLD snooping policy file filter on VLAN ports.

Syntax Description

<i>vlan_name</i>	Specifies a VLAN name.
<i>port_list</i>	Specifies one or more ports or slots and ports. On a SummitStack, can be a list of slots and ports. On a standalone switch, can be one or more port numbers. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.
<i>policy</i>	Specifies the policy file for the filter.

Default

None.

Usage Guidelines

Use this command to filter multicast groups to the specified VLAN ports.

The policy file used by this command is a text file that contains the IPv6 multicast addresses of the multicast groups that you wish to block.

To remove MLD snooping filtering from a port, use the **none** keyword version of the command.

Use the following template to create a snooping filter policy file:

```
#
# Add your group addresses between "Start" and "end"
# Do not touch the rest of the file!!!!
entry mldFilter {
  if match any {
    #----- Start of group addresses -----
    nlri FF03::1/128;
    nlri FF05::1/112;
    #----- end of group addresses -----
  } then {
    deny;
  }
  entry catch_all {
    if {
    } then {
    }
    permit;
  }
}
```

Example

The following example configures the policy file ap_multicast to filter multicast packets forwarded to VLAN marketing on ports 2:1-2:4:

```
configure mld snooping marketing ports 2:1-2:4 filter ap_multicast
```

History

This command was first available in ExtremeXOS 15.2.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MLD snooping feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure mld snooping vlan ports join-limit

```
configure mld snooping {vlan} vlan_name ports port_list join-limit
[num_joins | no-limit]
```

Description

Configures VLAN ports to support a maximum number of MLD joins.

Syntax Description

<i>vlan_name</i>	Specifies a VLAN name
<i>port_list</i>	Specifies one or more ports or slots and ports.
<i>num</i>	Specifies the maximum number of joins permitted on the ports. The range is 1 to 5000.

Default

No limit.

Usage Guidelines

None.

Example

The following example configures port 2:1 in the Default VLAN to support a maximum of 100 MLD joins:

```
configure mld snooping "Default" ports 2:1 join-limit 100
```

History

This command was first available in ExtremeXOS 15.2.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MLD snooping feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure mld ssm-map add

```
configure mld ssm-map add v6groupnetmask [v6sourceip | src_domain_name]
{ {vr} vr_name }
```

Description

Adds an MLD SSM Mapping entry on a VR.

Syntax Description

<i>v6groupnetmask</i>	You must provide group address with the mask length. Instead of configuring separate entries for a continuous range of IP addresses, this optimizes a range of group IP addresses to be configured as a single entry.
<i>v6sourceip</i>	Specifies the source IP address for which the SSM should apply.
<i>src_domain_name</i>	Provides the option to use DNS to obtain IP addresses dynamically by specifying the domain name.
VR <i>vr_name</i>	Specifies the virtual router name.

Default

N/A.

Usage Guidelines

When an MLDv1 report is received for this group or group range, the list of sources configured using this command is used as part of source-specific information to PIM.

The following error message displays when more than 50 source addresses are configured for a specific group:

```
ERROR: Cannot configure more than 50 sources for group ff30::1/128 on VR-Default
```

The following error message displays when a source address is already configured:

```
ERROR: Source 2001:0DB8:1::1 already present for group ff30::1/128 on VR-Default
```

The following error message displays when a DNS name is already configured:

```
ERROR: Only one source domain name allowed for group ff30::1/128 on VR-Default
```

Example

The following example configures a MLD-SSM mapping entry:

```
configure mld ssm-map add ff06::/64 2001::1
```

History

This command was first available in ExtremeXOS 15.5.

Platform Availability

This command is available on the platforms listed for the IPv6 multicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure mld ssm-map delete

```
configure mld ssm-map delete v6groupnetmask [v6sourceip |  
src_domain_name | all] [{vr} vr_name]
```

Description

Deletes an MLD SSM Mapping entry on a VR.

Syntax Description

<i>v6groupnetmask</i>	You must provide group address with the mask length. Instead of configuring separate entries for a continuous range of IP addresses, this optimizes a range of group IP addresses to be configured as a single entry.
<i>v6sourceip</i>	Specifies the source IP address for which the SSM should apply.
<i>src_domain_name</i>	Provides the option to use DNS to obtain IP addresses dynamically by specifying the domain name.
all	Specifies that all the mapping entries associated with <i>v6groupnetmask</i> are deleted.
vr <i>vr_name</i>	Specifies the virtual router name.

Default

N/A.

Usage Guidelines

When an MLDv1 report is received for this group or group range, the list of sources configured using this command is used as part of source-specific information to PIM.

The following error message displays when specified entry is not found:

```
ERROR: SSM Mapping entry (ff30::1/128, 2001:0DB8:1::10) not found on VR-Default
```

Example

The following example deletes a MLD-SSM mapping entry:

```
configure mld ssm-map delete ff06::/64 2001::1
```

History

This command was first available in ExtremeXOS 15.5.

Platform Availability

This command is available on the platforms listed for the IPv6 multicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure mpls add vlan

```
configure mpls add {vlan} vlan_name
```

Description

Adds an MPLS interface to the specified VLAN.

Syntax Description

<i>vlan_name</i>	Identifies the VLAN where the MPLS interface is added.
------------------	--

Default

VLANs are not configured with an MPLS interface.

Usage Guidelines

An MPLS interface must be configured on a VLAN in order to transmit or receive MPLS packets on that interface. By default, MPLS, LDP, and RSVP-TE are disabled for the MPLS interface. The specified VLAN should have an IP address configured and should have IP forwarding enabled. The MPLS interface on the VLAN does not become active until these two conditions are met. Also, if the IP address is unconfigured from the VLAN or IP forwarding is disabled for the VLAN, the MPLS interface goes down. The MPLS interface state is viewed using the `show mpls interface` command.

The VLAN must be operational for the MPLS interface to be up. This means that at least one port in the VLAN must be active or the VLAN must be enabled for loopback mode.

It is recommended that when you configure MPLS on an OSPF interface that can be used to reach a given destination, you should configure MPLS on all OSPF interfaces that can be used to reach that destination. (You should enable MPLS on all of the VLANs connected to the backbone network).

Example

The following example adds MPLS to the VLAN `vlan_usa`:

```
configure mpls add vlan vlan_usa
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure mpls delete vlan

```
configure mpls delete [{vlan} vlan_name | vlan all]
```

Description

Removes an *MPLS* interface from the specified *VLAN*.

Syntax Description

<i>vlan_name</i>	Identifies the VLAN for which the MPLS interface is deleted.
vlan all	Deletes the MPLS interface from all VLANS that have MPLS configured.

Default

VLANs are not configured with an MPLS interface.

Usage Guidelines

An MPLS interface must be configured on a VLAN in order to transmit or receive MPLS packets on that interface. If the MPLS interface is deleted, all configuration information associated with the MPLS interface is lost. Issuing this command brings down all LDP neighbor sessions and all LSPs that are established through the specified VLAN interface. When the all VLANs option is selected, the MPLS interface for all MPLS configured VLANs is deleted.

Example

The following example deletes MPLS from the VLAN *vlan_k*:

```
configure mpls delete vlan vlan_k
```

History

This command was first available in ExtremeXOS 11.6

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure mpls exp examination

```
configure mpls exp examination {value} value {qosprofile} qosprofile
```

Description

Configures the *QoS* profile that is used for the EXP value when EXP examination is enabled.

Syntax Description

<i>value</i>	Specifies the value that is used for the EXP value.
<i>qosprofile</i>	Specifies the QoS profile that is used for the EXP value.

Default

The QoS profile matches the EXP value + 1.

Usage Guidelines

This command configures the QoS profile that is used for the EXP value when EXP examination is enabled. By default, the QoS profile matches the EXP value + 1. That is, EXP value of 0 is mapped to QoS profile qp1, EXP value of 1 is mapped to QoS profile qp2, etc. This configuration has switch-wide significance. The EXP value must be a valid number from 0 through 7 and the qosprofile must match one of the switch's QoS profiles.



Note

EXP examination must be enabled using the “enable mpls exp examination” command before the configured EXP value to QoS profile mapping is actually used to process packets.

Example

The following command sets QoS profile q5 to be used for EXP value 7:

```
configure mpls exp examination value 7 qosprofile 5
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure mpls exp replacement

```
configure mpls exp replacement {qosprofile} qosprofile {value} value
```

Description

Configures the EXP value that is used for the specified [QoS](#) profile when EXP replacement is enabled.

Syntax Description

<i>qosprofile</i>	Specifies the QoS profile that is used for the EXP value.
<i>value</i>	Specifies the value that is used for the EXP value.

Default

The EXP value matches the QoS profile -1.

Usage Guidelines

This command configures the EXP value that is used for the QoS profile when EXP replacement is enabled. By default, the EXP value matches the QoS profile - 1. That is, QoS profile qp1 is mapped to EXP value of 0, QoS profile qp2 is mapped to EXP value of 1, etc. This configuration has switch-wide significance. The qosprofile must match one of the switch's QoS profiles and the EXP value must be a valid number from 0 through 7.



Note

EXP replacement must be enabled using the “enable mpls exp replacement” command before the configured EXP value to QoS profile mapping is actually used to process packets.

Example

The following command sets EXP value 2 to be used with QoS profile 4:

```
configure mpls exp replacement qosprofile qp4 value 2
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support *MPLS* as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure mpls labels max-static

```
configure mpls labels max-static max_static_labels
```

Description

Configures the number of labels that are reserved for specifying the incoming label for static LSPs and static pseudowires.

Syntax Description

labels	Specifies that labels are reserved to specify the incoming label for static LSPs and static pseudowires.
max-static	Specifies the number of labels that are reserved to specify the incoming label for static LSPs and static PWs.
<i>max_static-labels</i>	Specifies the value for the maximum number of static labels.

Default

The default static label range size is 100.

Usage Guidelines

Use this command to configure the number of labels that are reserved for specifying the incoming label for static LSPs and static PWs. The static label range generally starts at 16 and the default static label range size is 100. This means that the default static label range is 16 through 115 and can be allocated for either incoming (both transit and egress) static LSPs, or incoming static PWs. The maximum static label_range_size is equal to the incoming label table size – 100 labels for signaling. 960 labels are reserved for L3VPNs. The maximum number of labels available for static configuration is 7116, since at least 100 of those labels are reserved for dynamic signaling.

Since these values vary per-platform, use the `show mpls label usage` command to see details about label usage and platform capability. The minimum static label range size is 0.



Note

MPLS must be disabled when issuing this command. If MPLS is enabled, an error message is displayed and the command has no affect. All other labels, including outgoing labels for static LSPs and PWs and signaled labels used by RSVP-TE and LDP, are allocated out of the dynamic label space.

Example

The following example illustrates how to configure MPLS max-static labels, and how to display them:

```
Summit1.2 # show mpls lab usage

Label Type           Size      Label Range
-----
Supported            1048576   0x00000 - 0xfffff (0 - 1048575)
Reserved              16         0x00000 - 0x0000f (0 - 15)
Static                100        0x00010 - 0x00073 (16 - 115)
L3VPN                 960        0x00074 - 0x00433 (116 - 1075)
Dynamic              7116        0x00434 - 0x01fff (1076 - 8191)
Internal Use          0          0x00000 - 0x00000 (0 - 0)
...

Summit1.3 # disable mpls
* Summit1.4 # conf mpls lab max-static 7117
Error: There must be at least 100 dynamic labels remaining for MPLS signalling protocols.
* Summit1.5 # conf mpls lab max-static 7116
* Summit1.6 # show mpls lab usage
```

Label Type	Size	Label Range
-----	-----	-----
Supported	1048576	0x00000 - 0xfffff (0 - 1048575)
Reserved	16	0x00000 - 0x0000f (0 - 15)
Static	7116	0x00010 - 0x01bdb (16 - 7131)
L3VPN	960	0x01bdc - 0x01f9b (7132 - 8091)
Dynamic	100	0x01f9c - 0x01fff (8092 - 8191)
Internal Use	0	0x00000 - 0x00000 (0 - 0)
...		

History

This command was first available in ExtremeXOS 15.4

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure mpls ldp advertise

```
configure mpls ldp advertise [{direct [all | lsr-id | none]} | {rip [all
| none] | {static [all | none]}
```

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 <i>RIP</i> routes exported by <i>OSPF</i> .
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.
lsr-id	Specifies that an unsolicited label advertisement is originated for a direct route that matches the <i>MPLS</i> LSR ID.
none	Specifies that no unsolicited label mapping advertisements are originated for the specified route type.

Default

None—the default setting for RIP and static routing methods.

lsr-id—the default setting for direct routes.

Usage Guidelines

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**—Label mappings are originated for all routes of the specified type.
- **none**—No label mappings are originated for all routes of the specified type. This is the default setting for RIP and static routes.
- **lsr-id**—A label mapping is originated for a direct route that matches the MPLS LSR ID. This is the default setting for direct routes.

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 LDP to originate labels for all local IP interfaces:

```
configure mpls ldp advertise direct all
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure mpls ldp loop-detection

```
configure mpls ldp loop-detection [{hop-count hop_count_limit} {path-  
vector path_vector_limit}]
```

Description

Configures the loop-detection parameters used by LDP.

Syntax Description

hop-count	Configures the number of LSRs that the label message can traverse.
<i>hop_count_limit</i>	Specifies the hop count limit. The valid configuration range is from 1 to 255.
path-vector	Configures the maximum number of LSR IDs that can be propagated in the label message.
<i>path_vector_limit</i>	Specifies the path vector limit. The valid configuration range is from 1 to 255.

Default

The default for the hop-count and path-vector limits is 255.

Usage Guidelines

Configuration changes are only applicable to newly created LDP sessions. Disabling and enabling LDP forces all the LDP sessions to be recreated. LDP loop detection must first be enabled for these configuration values to be used.

Example

This command sets the LDP hop count loop detection value to 10. The configured path vector value remains at 255.

```
configure mpls ldp loop-detection hop-count 10
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure mpls ldp timers

```
configure mpls ldp timers [targeted | link] [{hello-time  
    hello_hold_seconds} {keep-alive-time keep_alive_hold_seconds}]
```

Description

Configures LDP peer session timers for the switch.

Syntax Description

targeted	Specifies targeted LDP sessions.
link	Specifies link LDP sessions.
<i>hello_hold_seconds</i>	The amount of time (in seconds) that a hello message received from a neighboring LSR remains valid. The rate at which Hello messages are sent is 1/3 the configured hello-time. If a Hello message is not received from a particular neighboring LSR within the specified <i>hello_hold_seconds</i> , then the hello-adjacency is not maintained with that neighboring LSR. The range is 6 to 65,534 seconds.
<i>keep_alive_hold_seconds</i>	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 <i>keep_alive_hold_seconds</i> , the corresponding LDP session is torn down. The range is 6 to 65,534 seconds.

Default

link *hello_hold_seconds* - 15 seconds

targeted *hello_hold_seconds* - 45 seconds

link *keep_alive_hold_seconds* - 40 seconds

targeted *keep_alive_hold_seconds* - 60 seconds

Usage Guidelines

The LDP peer hello-adjacency timers are separately configurable for link and targeted LDP sessions. The hello timer parameter specifies the amount of time (in seconds) that a Hello message received from a neighboring LSR remains valid. The rate at which Hello messages are sent is 1/3 the configured hello-time. If a Hello message is not received from a particular neighboring LSR within the specified *hello_hold_seconds*, then the hello-adjacency is not maintained with that neighboring LSR.

The session *keep_alive_hold_seconds* parameter specifies the time (in seconds) during which an LDP message must be received for the LDP session to be maintained. The rate at which Keep Alive messages are sent, provided there are no LDP messages transmitted, is 1/6 the configured keep-alive-time. If an LDP PDU is not received within the specified session *keep_alive_hold_seconds* interval, the corresponding LDP session is torn down. The minimum and maximum values for hold timers are 6 and 65,534, respectively.

Changes to targeted timers only affect newly created targeted sessions. Disabling and then enabling VPLS or LDP causes all current targeted sessions to be re-created. The default values for the various times are as follows: link *hello_hold_seconds* (15), link *keep_alive_hold_seconds* (40), targeted *hello_hold_seconds* (45), and targeted *keep_alive_hold_seconds* (60). Changes to the link keep-alive timers do not take effect until the LDP session is cycled.

Example

The following command configures link-level LDP hello adjacency hold time to 30 seconds and the keep alive time to 10 seconds:

```
configure mpls ldp timers link hello-time 30 keep-alive-time 10
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure mpls lsr-id

```
configure mpls lsr-id ipaddress
```

Description

Configures the [MPLS](#) LSR ID for the switch.

Syntax Description

<i>ipaddress</i>	Specifies an IP address to identify the MPLS LSR for the switch. The MPLS LSR-ID should be configured to the same IP address as the OSPF Router ID.
------------------	---

Default

No LSR ID is configured by default.

Usage Guidelines

LDP, RSVP-TE, and L2 VPNs all use the LSR ID. It is normally set to the OSPF Router ID.

The LSR ID must be configured before MPLS can be enabled. The LSR ID cannot be changed while MPLS is enabled. It is highly recommended that an IP address be configured on a OSPF enabled loopback [VLAN](#) that matches the configured LSR ID and OSPF ID. If an LSR ID loopback IP address is configured, OSPF automatically advertises the LSR ID as a routable destination for setting up LSPs. The LSR ID remains active if an interface goes down if the LSR-ID is configured as an IP address on a loopback VLAN, as recommended. This significantly enhances network stability and operation of an MPLS network.

Example

The following command configures the LSR ID to 192.168.50.5:

```
configure mpls lsr-id 192.168.50.5
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure mpls rsvp-te bandwidth committed-rate

```
configure mpls rsvp-te bandwidth committed-rate committed_bps [Kbps |  
  Mbps | Gbps] [{vlan} vlan_name | vlan all] {receive | transmit |  
  both}
```

Description

Specifies the maximum amount of Committed Information Rate (CIR) bandwidth which can be used by RSVP-TE LSP reservations.

Syntax Description

<i>committed_bps</i>	Specifies a bitrate for the bandwidth to be reserved.
Kbps	Specifies the designated bitrate in kilobits per second.
Mbps	Specifies the designated bitrate in megabits per second.
Gbps	Specifies the designated bitrate in gigabits per second.
vlan	Specifies that the bandwidth is to be reserved for a specific VLAN .
<i>vlan_name</i>	Identifies the VLAN for which the bandwidth is reserved.
vlan all	Specifies that the bandwidth is reserved for all VLANs that have MPLS configured.
receive	Specifies that the bandwidth is reserved for ingress traffic only.
transmit	Specifies that the bandwidth is reserved for egress traffic only.
both	Specifies that the bandwidth is reserved for both ingress and egress traffic.

Default

The default is zero, which means no RSVP-TE LSP bandwidth reservations are accepted.

If bandwidth is specified without specifying traffic direction, the default is both directions.

Usage Guidelines

This command specifies the maximum amount of Committed Information Rate (CIR) bandwidth which can be used by dynamic RSVP-TE LSP bandwidth reservations. By sub-allocating reserveable bandwidth for RSVP-TE from the VLAN's available bandwidth, the switch can guarantee that as LSPs are established, a minimum amount of CIR bandwidth is available for other traffic.



Note

Beginning with ExtremeXOS Release 12.2.1, CIR bandwidth for the receive direction is not tracked by TE IGPs, such as *OSPF-TE*, and configuring it is not required. Configuring CIR bandwidth for the receive direction does not prevent an LSP from going operational due to lack of receive bandwidth; however, it can be useful for tracking and informational purposes. An Info level log (MPLS.RSVPTE.IfRxBwdthExcd) is generated if the setup of a TE LSP requires receive bandwidth greater than that which is currently available for the receive direction on a particular interface. This generally happens only when TE LSPs with different previous hops ingress the switch on the same interface (for example, from a multi-access link) and egress the switch on different interfaces.

The keyword **both** configures the reserved bandwidth for both ingress and egress LSP CIR reservations and overwrites any previous receive or transmit settings.

Example

The following command reserves 25 Mbps of CIR bandwidth for all RSVP-TE CIR reservations on the specified VLAN:

```
configure mpls rsvp-te bandwidth committed-rate 25 Mbps vlan vlan_10
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure mpls rsvp-te lsp add path

```
configure mpls rsvp-te lsp lsp_name add path [path_name | any] {profile profile_name} {primary {frr_profile_name} | secondary}
```

Description

Adds a configured path to the specified RSVP-TE LSP.

Syntax Description

<i>lsp_name</i>	Specifies the name of the LSP you are configuring.
<i>path_name</i>	Specifies the name of the path to be used by the specified LSP.
any	Configures the specified LSP to use any path.
<i>profile_name</i>	Specifies a profile to be applied to the specified LSP. If the profile name is omitted, the profile named default is used.
primary	Designates the specified path as the primary path. Only one primary path can be configured for an RSVP-TE LSP. If this option is omitted and no primary path has been specified, the specified path is added as a primary path. If not specified and a primary path has already been added, the path is added as a secondary path.
secondary	Designates the specified path as a secondary path.
<i>frr_profile_name</i>	Specifies a fast reroute (FRR) profile to be applied to the detour LSP that backs up the specified LSP.

Default

N/A.

Usage Guidelines

The LSP is not signaled until a path is added to the LSP.

If you want fast reroute protection for the LSP, use the primary option and specify the fast reroute profile name you want to use. To specify the default fast reroute profile, enter default-frr.

The switch chooses the local *MPLS VLAN* interface from which to signal the LSP. To force an LSP to use a specific local MPLS interface, configure the local interface IP address as the first ERO in the associated path.

Example

This command adds the path sydney-bypass to the LSP named aus as a secondary path:

```
configure mpls rsvp-te lsp aus add path sydney-bypass secondary
```

History

This command was first available in ExtremeXOS 11.6.

The fast reroute capability was added in ExtremeXOS 12.1.

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure mpls rsvp-te lsp change

```
configure mpls rsvp-te lsp lsp_name change [path_name | any] use profile
[{standard_profile_name} {frr_profile_name}]
```

Description

Changes the configuration that has been configured with the `configure mpls rsvp-te lsp lsp_name add path [path_name | any] {profile}profile_name {primary} {frr_profile_name} | secondary}` command.

Syntax Description

<i>lsp_name</i>	Specifies the name of the LSP you are changing.
<i>path_name</i>	Specifies the name of the path to be used by the specified LSP.
any	Configures the specified LSP to use any path.
<i>standard_profile_name</i>	Specifies a profile to be applied to the specified LSP. If the profile name is omitted, the profile named default is used.
<i>frr_profile_name</i>	Specifies a fast reroute (FRR) profile to be applied to the detour LSP that backs up the specified LSP.

Default

N/A.

Usage Guidelines

None.

Example

This command changes the LSP named aus to use any available path:

```
configure mpls rsvp-te lsp aus change any
```

History

This command was first available in ExtremeXOS 11.6.

The fast reroute capability was added in ExtremeXOS 12.1.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure mpls rsvp-te lsp delete path

```
configure mpls rsvp-te lsp lsp_name delete path [path_name | any | all]
```

Description

Deletes a path from the specified RSVP-TE LSP.

Syntax Description

<i>lsp_name</i>	Specifies a name for the RSVP-TE LSP.
<i>path_name</i>	Specifies a name for the path to be deleted from the RSVP-TE LSP.
any	Configures the specified LSP to use any path.
all	Deletes all added paths from the specified RSVP-TE LSP.

Default

N/A.

Usage Guidelines

This command deletes a path from the specified RSVP-TE LSP. All the added paths can be deleted by specifying the **all** keyword. If the active path is deleted, then one of the other configured paths becomes the active path for the LSP. If there are no other defined paths, then the LSP is marked down and cannot be used to forward IP or VPN traffic.

Example

The following command deletes the path called through-knightsbridge for the LSP london:

```
configure mpls rsvp-te lsp london delete path through-knightsbridge
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure mpls rsvp-te lsp fast-reroute

```
configure mpls rsvp-te lsp lsp_name fast-reroute [enable | disable]
```

Description

Enables or disables fast-reroute protection for the specified LSP.

Syntax Description

<i>lsp_name</i>	Specifies the name of the LSP you are configuring.
-----------------	--

Default

Disabled.

Usage Guidelines

To signal the fast-reroute protected LSP, use the `enable mpls rsvp-te lsp [lsp_name | all]` command. Similarly, to disable the fast-reroute protected LSP, use the `disable mpls rsvp-te lsp [lsp_name | all]` command.

Example

This command enables fast-reroute protection on LSP aus:

```
configure mpls rsvp-te lsp aus fast-reroute enable
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure mpls rsvp-te lsp path use profile

```
configure mpls rsvp-te lsp lsp_name path [path_name | any] use profile
profile_name
```

Description

Changes the profile that the configured LSP path uses.

Syntax Description

<i>lsp_name</i>	Specifies the RSVP-TE LSP.
<i>path_name</i>	Specifies the configured RSVP-TE LSP path.
<i>profile_name</i>	Specifies a profile to be applied to the configured LSP path.

Default

N/A.

Usage Guidelines

This command changes the profile that the configured LSP path uses.



Note

Changing the profile while an LSP is active may cause the LSP to be torn down and re-signaled.

Example

The following command configures the switch to apply the LSP profile gold-class to the LSP path sydney-bypass for the LSP aus:

```
configure mpls rsvp-te lsp aus path sydney-bypass use profile gold-class
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure mpls rsvp-te lsp transport

```
configure mpls rsvp-te lsp lsp_name transport [ip-traffic [allow | deny]  
| vpn-traffic [allow {all | assigned-only} | deny]]
```

Description

Configures the type of traffic that may be transported across a named LSP.

Syntax Description

<i>lsp_name</i>	Specifies the RSVP-TE LSP.
ip-traffic	Controls the forwarding of routed IP traffic across the specified LSP.
vpn-traffic	Controls the forwarding of VPN traffic over the LSP.
allow	Allows transport of the specified traffic across the LSP.
deny	Denies transport of the specified traffic across the LSP.
allow	Allows all VPLS VPN traffic to be transported across the LSP.
all	Allows the transmission of all VPN traffic over the LSP.
assigned-only	Limits the transport of VPN traffic to VPLS instances that are explicitly configured to use the specified LSP name.

Default

The default behavior is to allow RSVP-TE LSPs to transport all types of traffic without restriction.

Usage Guidelines

This command configures the type of traffic that may be transported across a named LSP. By default, both IP traffic and VPN traffic are set to allow transport for a newly created LSP. The `ip-traffic` keyword is used to allow or deny forwarding of routed IP traffic across the specified LSP. If allowed, the LSP label information is inserted into the routing table and the switch forwards traffic over the LSP that matches the IP route entry to which this LSP is associated. If denied, the LSP label information is removed from the routing table and the switch does not use the LSP to transport IP traffic. The `vpn-traffic` keyword controls the transmission of VPN traffic over the LSP. When denied, the LSP is not used as a transport for PWs or other VPN related traffic. These transport configuration options are independent. For example, if `vpn-traffic` is set to allow and `ip-traffic` is set to deny, then no routed IP traffic is transported across the LSP, but the LSP may still be used to transport VPN traffic.

The optional `assigned-only` keyword limits the transport of VPN traffic to only those VPLS instances that are explicitly configured to use the specified LSP name.

Example

The following command prevents the switch from using LSP `aus` to forward IP traffic:

```
configure mpls rsvp-te lsp aus transport ip-traffic deny
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support *MPLS* as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure mpls rsvp-te metric

```
configure mpls rsvp-te metric [value | use-igp] {vlan} vlan_name
```

Description

Configures the TE metric value for the RSVP-TE interface specified by the *vlan_name* argument.

Syntax Description

<i>value</i>	Specifies a value for the RSVP-TE metric.
vlan	Specifies that the RSVP-TE metric is configured for a specific VLAN .
<i>vlan_name</i>	Identifies the VLAN for which the RSVP-TE metric is configured.

Default

The associated default IGP metric.

Usage Guidelines

The TE metric can be any unsigned non-zero 32-bit integer. The default value for the RSVP-TE interface is to use the associated default IGP metric. The TE metric is exchanged between [OSPF](#) routers and is used in the calculation of the CSPF topology graph.

Example

The following command configures an RSVP-TE metric of 220 on the specified VLAN:

```
configure mpls rsvp-te metric 220 vlan vlan_10
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure mpls rsvp-te path add ero

```
configure mpls rsvp-te path path_name add ero [ { include } ipNetmask  
[strict|loose] | exclude ipNetmask] {order number}
```


Description

The routed path for an RSVP-TE LSP can 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. LSRs and/or subnets can be either included or excluded.

Syntax Description

<i>path_name</i>	Specifies the path to which the IP address is added.
include	Specifies an LSR or subnet to be included in the path calculation.
<i>ipNetmask</i>	Specifies an IP prefix.
strict	Specifies that the subobject must be topologically adjacent to the previous subobject in the ERO list.
loose	Specifies that the subobject need not be topologically adjacent to the previous subobject in the ERO list.
exclude	Specifies a subnet to be excluded in the path calculation.
<i>number</i>	Specifies the LSR path order.

Default

The order value defaults to 100 if the path has no EROs configured or a value 100 more than the highest order number configured for the path.

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 that the path traverses. Each defined LSR or subnet represents an ERO subobject. Up to 64 subobjects can be added to each path name. The ERO keyword identifies an LSR using an IP prefix, which may represent an LSR's Router ID, loopback address, or direct router interface. Each IP prefix is included in the ERO as an IPv4 subobject.

If the ERO is specified as strict, the strict subobject must be topologically adjacent to the previous subobject as listed in the ERO. If the ERO is specified as loose, the loose subobject is not required to be topologically adjacent to the previous subobject as listed in the ERO. If the specified IP prefix matches the OSPF router ID or a configured loopback IP address, the ERO must be configured as loose.

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. Thus, the LSP path follows the configured path of IP prefixes with a number value from low to high. If the order keyword is not specified, the number value for the LSR defaults to a value equal to the current highest number value plus 100. 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.

¹ "Topologically adjacent" indicates that the router next hop matches either the interface IP address or [*OSPF*](#) router ID of an immediate peer LSR.

The order of a configured subobject cannot be changed. The ERO subobject must be deleted and re-added with 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 re-signaled using the new ERO. Duplicate ERO subobjects are not allowed.

Defining an ERO for the path is optional. If no ERO is configured, the path is signaled along the best available 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 available path. If the next subobject in the ERO is loose, the best available path to the next subobject is chosen. Configuring EROs could lead an LSP to take an undesirable path through the network, so care should be taken when specifying EROs.

Example

The following example adds the IP interface address 197.57.30.7/24 as a loose ERO to the path sydney-bypass:

```
configure mpls rsvp-te path sydney-bypass add ero 197.57.30.7/24 loose
```

History

This command was first available in ExtremeXOS 11.6.

The **include** and **exclude** options were added in ExtremeXOS 15.7. "Include" was the previous default behavior.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure mpls rsvp-te path delete ero

```
configure mpls rsvp-te path path_name delete ero [all | ipNetmask | order number]
```

Description

Deletes a subobject from the Explicit Route Object (ERO) for the specified path name.

Syntax Description

<i>path_name</i>	Specifies the path from which the ERO is deleted.
all	Specifies that the entire ERO should be deleted from the named path.
<i>ipNetmask</i>	Specifies the ERO subobject to be deleted.
<i>number</i>	Specifies the order number of the ERO subobject to be deleted.

Default

N/A.

Usage Guidelines

This command deletes a subobject from the Explicit Route Object (ERO) for the specified path name. The ERO subobject is specified using an IP prefix or order number. If a subobject is deleted from an ERO while the associated LSP is established, the path is torn down and is re-signaled using a new ERO. The all keyword may be used 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 available path and no ERO is included in the path message.

Example

The following command deletes all the configured EROs from the path sydney-bypass:

```
configure mpls rsvp-te path sydney-bypass delete ero all
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure mpls rsvp-te profile (fast-reroute)

```
configure mpls rsvp-te profile frr_profile_name {bandwidth
  bandwidth_rate_bps bandwidth_rate_unit} {detour {hop-limit
  hop_limit_value} {bandwidth-protection [enabled | disabled]} {node-
protection [enabled | disabled]}} {hold-priority hold_priority_value}
  {setup-priority setup_priority_value}
```

Description

Configures the specified RSVP-TE FRR profile.

Syntax Description

<i>frr_profile_name</i>	Specifies the FRR LSP profile to configure.
<i>bandwidth_rate_bps</i>	Specifies the bandwidth requirement for the FRR LSP. This should be set to match the options chosen for the protected LSP. Otherwise, a mismatch between the bandwidth settings for the detour and protected LSPs can impact service.

<i>bandwidth_rate_unit</i>	Specifies the units for the bandwidth rate. Valid entries are Kbps, Mbps, and Gbps.
detour	Specifies the detour method of fast reroute. This is the only method supported in this release.
<i>hop_limit_value</i>	Specifies the maximum number of hops that the detour path is allowed to take from the current node or point of local repair (PLR) to a merge point (MP) node. If set to 0, only link protection is provided.
bandwidth-protection	When enabled, this option specifies that the signaled bandwidth on the detour path must be guaranteed. If this option is disabled, the detour path might not support the bandwidth needed for the protected LSP.
node-protection	When enabled, the this option indicates to the PLRs along a protected path that a detour path that bypasses at least the next node of the protected LSP is desired. If this option is disabled, the backup path might or might not bypass the next node, in which case the user might or might not have next-node protection.
hold-priority	Specifies the hold priority of the LSP. Lower numbers indicate higher priority. The range is from 0 to 7. Hold priority is used when deciding whether a session can be preempted by another session. This works exactly the same as the hold-priority set in the standard profile that is valid for the protected LSP and for standard LSPs.
setup-priority	Specifies the setup priority of the LSP. Lower numbers indicate higher priority. The range is from 0 to 7. The setup priority is used when deciding whether the detour LSP can preempt another session. This works exactly the same as the setup-priority set in the standard profile that is valid for the protected LSP and standard LSPs.

Default

Bandwidth: Newly-created profiles are configured as best-effort. Setup-priority: 7 (lowest) Hold-priority: 0 (highest) Hop-limit: 3 Protect-bandwidth: enabled Protect-node: enabled

Usage Guidelines

A FRR profile is a set of attributes that are applied to the detour and protected LSPs when a protected LSP is configured. A default profile (frr-default) is provided which cannot be deleted, but can be applied to any protected LSP. The maximum number of configurable profiles is 1000.



Note

Changing any of the profile parameters causes LSPs using the profile to be torn down and re-signaled. There is no guarantee that the re-signaled LSP will be successfully established. Future ExtremeXOS implementations may support the make-before-break LSP concept.

Example

The following command configures the FRR profile frrprofile for 100 Mbps bandwidth:

```
configure mpls rsvp-te profile frrprofile bandwidth 100 Mbps
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure mpls rsvp-te profile

```
configure mpls rsvp-te profile profile_name {bandwidth [best-effort |
  [{committed-rate committed_bps [Kbps | Mbps | Gbps]} {max-burst-size
  burst_size [Kb | Mb]} {peak-rate peak_bps [Kbps | Mbps | Gbps]}]}
  {hold-priority hold_priority} {mtu [number | use-local-interface]}
  {path-computation [full | partial]} {record [enabled {route-only} |
  disabled]} {setup-priority setup_priority}
```

Description

Configures an RSVP-TE profile with the specified profile name.

Syntax Description

<i>profile_name</i>	Specifies the LSP profile.
bandwidth	Specifies bandwidth reservation.
best-effort	Indicates no bandwidth reservation.
<i>committed_bps</i>	Specifies the committed bandwidth to be reserved across the MPLS network, in bits per second. The range is from 64 Kbps to 10 Gbps.
<i>peak_bps</i>	Specifies the maximum bandwidth signaled in bits per second. The range is from 64 Kbps to 10 Gbps.
Kbps	Specifies the designated bitrate in kilobits per second.
Mbps	Specifies the designated bitrate in megabits per second.
Gbps	Specifies the designated bitrate in gigabits per second.
<i>burst_size</i>	Specifies the maximum number of bytes (specified in bits) that the LSP is allowed to burst above the specified peak-rate. The range is from 0 to 1000 Mb.
Kb	Kilobits
Mb	Megabits
<i>hold_priority</i>	Specifies the priority of the LSP. Lower numbers indicate higher priority. The range is from 0 to 7.
<i>setup_priority</i>	Specifies the priority of the LSP. Lower numbers indicate higher priority. The range is from 0 to 7.
<i>number</i>	Specifies the MTU value for the LSP. The range is from 296 to 9216/

use-local-interface	Specifies that the MTU value is inherited from the local egress VLAN interface.
record	Configures hop-by-hop path recording.
enabled route-only	Causes the Record Route Object (RRO) to be inserted into the path message. The enabled option enables recording of hops and labels. The enabled route-only option records only hops.
disabled	Specifies that no RRO is inserted into the path message.
path-computation	Computation strategy for calculating a path to the LSP destination: <ul style="list-style-type: none"> full = Requires the ingress node to fully calculate a path to the LSP destination (default). partial = Allows the ingress node to calculate only part of the path to the LSP destination.
full	Allows the entire LSP path to be specified at ingress LSR (no calculations performed by any transit nodes).
partial	Allows you to specify 'part' of the path at ingress LSR. (For OSPF usage, specify LSP path to the ABR, then ABR provides the calculation into the other areas.)

Default

Bandwidth: Newly-created profiles are configured as best-effort.

Setup-priority: 7 (lowest).

Hold-priority: 0 (highest).

Path recording: disabled.

MTU: use-local-interface.

Path-computation: full.

Usage Guidelines

A profile is a set of attributes that are applied to the LSP when the LSP is configured using the `configure mpls rsvp-te lsp` command. A default profile is provided which cannot be deleted, but may be applied to any TE 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 1000.

LSPs may signal reserved bandwidth. By default, newly created profiles are configured to not signal bandwidth requirements and thus are classified as best-effort. If bandwidth needs to be reserved across the MPLS network, the bandwidth parameters specify the desired reserved bandwidth for the LSP. The committed-rate specifies the mean bandwidth and the peak-rate specifies the maximum bandwidth signaled. The peak-rate must be equal to or greater than the committed-rate. If the peak-rate is not specified, traffic is not clipped above the committed-rate setting. The rates are specified in bps and must be qualified by Kbps, Mbps, or Gbps. The minimum and maximum bandwidth rates are 64 Kbps and 10 Gbps, respectively. The max-burst-size specifies the maximum number of bytes (specified in bits) that the LSP is allowed to burst above the specified peak-rate. The minimum burst size is 0 and the maximum burst size is 1000 Mb.

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 (lowest). The hold-priority range is also 0 to 7 and the default value is 0 (highest). If bandwidth is requested for the LSP, the CSPF calculation uses the available bandwidth associated with the CoS as specified by the hold-priority.

The bandwidth, hold-priority, and setup-priority values are signaled in the path message. If the bandwidth setting is changed, all LSPs using this profile are re-signaled. If the bps setting is decreased, a new path message is sent along the LSP indicating the new reservation. If the bps setting is increased, the LSP is torn down and resigaled using the new bandwidth reservations.

The record command 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 RESV 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.

The mtu keyword optionally specifies the MTU value for the LSP. By default, this value is set to use-local-interface. In the default configuration, the MTU value is inherited from the local egress VLAN interface. The minimum MTU value is 296 and the maximum value is 9216. Path MTU information is carried in the Integrated Services or Null Service RSVP objects and is used by RSVP to perform path MTU identification.



Note

Changing any of the profile parameters causes LSPs using the profile to be torn down and re-signaled. There is no guarantee that the re-signaled LSP will be successfully established. Future ExtremeXOS implementations may support the make-before-break LSP concept.

To view a profile configuration, enter the following command:

```
show mpls rsvp-te profile {profile_name} {detail}
```

To view LSP recorded route information, enter one of the following commands:

```
show mpls rsvp-te lsp [ingress {fast-reroute} | ingress_lsp_name
| ingress ingress_lsp_name | ingress [destination | origin] ipaddress]
[{all-paths | detail] | summary | down-paths {detail}] show mpls rsvp-
te lsp [egress | transit] {fast-reroute} [{lsp_name} [{destination |
origin] ipaddress} {detail} | summary}
```

Example

The following command configures the RSVP-TE profile gold-class with a committed bandwidth of 100 Mbps and the setup and hold priorities are both set to 0 (highest priority):

```
configure mpls rsvp-te profile gold-class bandwidth committed-rate 100 mbps hold-priority
0 setup-priority 0
```

History

This command was first available in ExtremeXOS 11.6.

The **path-computation** option added in ExtremeXOS 21.1

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure mpls rsvp-te timers lsp rapid-retry

```
configure mpls rsvp-te timers lsp rapid-retry {decay-rate percent}
      {delay-interval milliseconds} {retry-limit [number]}
```

Description

Configures the timers associated with rapidly retrying failed LSPs.

Syntax Description

<i>percent</i>	Specifies a percent increase in the interval allowed before each subsequent attempt to re-signal an LSP. The valid range is from 0 to 100 percent.
<i>milliseconds</i>	Specifies the time (in milliseconds) to wait before attempting to re-signal the LSP.
retry-limit	Specifies the maximum allowed attempts to establish an LSP.
<i>number</i>	Specifies a maximum number of allowed attempts to establish an LSP. The valid number range is from zero to 255.

Default

Delay interval: 500 milliseconds.

Decay rate: 50%.

Retry limit: 10.

Usage Guidelines

This command configures the timers associated with rapidly retrying failed LSPs. If an LSP fails to establish, the switch attempts to rapidly retry the setup by sending additional path messages based on the rapid-retry timers. The delay-interval timer specifies the time (in milliseconds) to wait before sending another path message. If the LSP fails to establish itself on subsequent attempts, the delay-interval time is incremented based on the decay-rate setting. The decay operation multiplies the delay-interval time by the decay rate, and adds the result to the current delay-interval time.

For example, if the decay-rate is set to 50 percent and the current delay-interval time is 500 milliseconds, a path message is retransmitted in 750 milliseconds. If the LSP fails to establish on the next attempt, a path message is retransmitted after a further decayed delay interval of 1125 milliseconds (1.125 seconds). A per-LSP delay-interval time is maintained for each LSP until the LSP is established. This process of decaying the retry time continues until the LSP is established or the retry-limit expires. If the retry-limit is reached, attempts to rapidly retry the LSP are suspended.

When the switch starts the process of re-signaling the LSP based on the standard-retry timers, the LSP's rapid-retry timers return to the initial configuration settings. If the standard-retry delay-interval time is reached before all of the rapid-retry attempts have completed, the standard-retry mechanisms take over.

The default rapid-retry LSP timer parameter values are 500 milliseconds for the delay-interval, 50 percent for the decay-rate, and a retry-limit of 10. The valid range for delay-interval is 10 to 1000 milliseconds. The valid decay-rate range is 0 to 100 percent. The valid retry-limit is 0 to 100. A value of 0 indicates that the LSP is not re-signaled using the rapid-retry timers.

When summary-refresh or bundle-message is enabled, the rapid-retry timer values are used for resending any message that is not acknowledged.

Example

The following command sets the maximum number of rapid retries to five:

```
configure mpls rsvp-te timers lsp rapid-retry retry-limit 5
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure mpls rsvp-te timers lsp standard-retry

```
configure mpls rsvp-te timers lsp standard-retry {decay-rate percent}  
        {delay-interval seconds} {retry-limit [number | unlimited]}
```

Description

Configures the timers associated with the establishment of an LSP.

Syntax Description

<i>percent</i>	Specifies a percent increase in the interval allowed before each subsequent attempt to re-signal an LSP. The valid range is from 0 to 100 percent.
<i>seconds</i>	Specifies the time (in seconds) to wait before attempting to re-signal the LSP.
retry-limit	Specifies the maximum allowed attempts to establish an LSP.

<i>number</i>	Specifies a maximum number of allowed attempts to establish an LSP. The valid number range is from zero to 255.
unlimited	Allows unlimited attempts to establish an LSP.

Default

- Delay interval: 30 seconds.
- Decay rate: 0%.
- Retry limit: unlimited.

Usage Guidelines

This command configures the timers associated with the establishment of an LSP. If an LSP fails to establish, the LSP is re-signaled based on the configuration of these timers. The delay-interval timer specifies the time (in seconds) to wait before attempting to re-signal the LSP. If the LSP fails to establish itself on subsequent attempts, the delay-interval time is incremented based on the decay-rate setting. The decay operation multiplies the delay-interval time by the decay rate, and adds the result to the current delay-interval time. For example, if the decay-rate is set to 50 percent and the current delay-interval time is 30 seconds, the LSP is re-signaled in 45 seconds. If the LSP failed to establish on the next attempt, the delay interval would be further decayed to 67 seconds.

A per-LSP delay-interval time is maintained for each LSP until the LSP is established. This operation of decaying the retry time continues until the LSP is established or the retry-limit expires. If the retry-limit is reached, attempts to establish the LSP are suspended.

Disabling and enabling the LSP resets the LSP's delay-interval time and retry-limit to the initial configuration settings and LSP establishment attempts resume. The default LSP timer parameter values are 30 seconds for delay-interval, with a 0 percent decay-rate, and retry-limit of unlimited. The valid range for delay-interval is 1 to 60 seconds. The valid decay-rate range is 0 to 100 percent. The valid retry-limit is 0 to 255 or unlimited. A value of 0 indicates that the LSP is not re-signaled.

Example

The following command allows unlimited retries for establishing MPLS RSVP-TE LSPs:

```
configure mpls rsvp-te timers lsp standard-retry retry-limit unlimited
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure mpls rsvp-te timers session

```
configure mpls rsvp-te timers session [{bundle-message-time bundle_message_milliseconds} {hello-keep-multiplier
hello_keep_number} {hello-time hello_interval_seconds} {refresh-keep-multiplier refresh_keep_number} {refresh-time refresh_seconds}
{summary-refresh-time summary_refresh_milliseconds}] [{vlan}
vlan_name | vlan all]
```

Description

Configures the RSVP-TE protocol parameters for the specified *VLAN*.

Syntax Description

<i>bundle_message_milliseconds</i>	Specifies the maximum time a transmit buffer is held to allow multiple RSVP messages to be bundled into a single PDU. The valid range is from 50 to 3000 milliseconds.
<i>hello_keep_number</i>	Specifies the number of hello-time intervals that can elapse before an RSVP-TE peer is declared unreachable. The range is from one to 255.
<i>hello_interval_seconds</i>	Specifies the RSVP Hello packet transmission interval. The valid range is from 1 to 60 seconds.
<i>refresh_keep_number</i>	Specifies a factor to be used in calculating the maximum allowed interval without an RSVP refresh message before an RSVP session is torn down. The range is from one to 255.
<i>refresh_seconds</i>	Specifies the interval for sending refresh path messages. The range is from 1 to 600 seconds.
<i>summary_refresh_milliseconds</i>	Specifies the interval for sending summary refresh messages. The valid range is from 50 (1/20 second) to 10000 (10 seconds).
vlan	Specifies that the configured protocol parameters are for a specific VLAN.
<i>vlan_name</i>	Identifies a particular VLAN for which the protocol parameters are configured.
vlan all	indicates that the protocol configuration parameters apply to all RSVP-TE enabled VLANs.

Default

Bundle-message-time: 1000 milliseconds (1 second).

Hello-keep-multiplier value: 3.

Hello-time: 3 seconds.

Refresh-keep-multiplier value: 3.

Refresh-time: 30 seconds.

Summary-refresh-time: 3000 milliseconds (3 seconds).

Usage Guidelines

This command configures the RSVP-TE protocol parameters for the specified VLAN. The VLAN keyword all indicates that the configuration changes apply to all VLANs that have been added to MPLS.

The hello-time value specifies the RSVP hello packet transmission interval. The RSVP hello packet enables the switch to detect when an RSVP-TE peer is no longer reachable. If an RSVP hello packet is not received from a peer within the configured interval, the peer is declared down and all RSVP sessions to and from that peer are torn down. The formula for calculating the maximum allowed interval is: $[\text{hello-time} * \text{hello-keep-multiplier}]$. The default hello-interval time is 3 seconds with a valid range from 1 to 60 seconds. The default hello-keep-multiplier value is three with a range from one to 255.

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 enable the switch to detect when an LSP is no longer active. Path messages are used to refresh the LSP if summary refresh is disabled. If summary refresh is enabled, summary refresh messages are sent in place of sending individual path messages for every LSP. The default refresh-time is 30 seconds. The minimum and maximum refresh-time values are one and 600 (or 10 minutes) respectively.

If summary refresh is enabled, summary refresh messages are sent at intervals represented by the configured summary-refresh-time. The configurable summary-refresh-time range is 50 milliseconds (one twentieth of a second) to 10000 milliseconds (10 seconds). The default setting for summary-refresh-time is 3000 milliseconds (3 seconds). RSVP sessions are torn down if an RSVP refresh message is not received from a peer within the configured interval. The formula for calculating the maximum allowed interval is: $[(\text{refresh-keep-multiplier} + 0.5) * 1.5 * (\text{refresh-time or summary-refresh-time})]$. The default refresh-keep-multiplier value is three. The minimum and maximum refresh-keep-multiplier values are one and 255 respectively.

The bundle-message-time, specified in milliseconds, indicates the maximum time a transmit buffer is held to allow multiple RSVP messages to be bundled into a single PDU. The default bundle-message-time is 1000 milliseconds (one second). The bundle-message-time value may be set to any value between 50 milliseconds and 3000 milliseconds (or 3 seconds). Message bundling is only attempted when it is enabled.



Note

Summary refresh must be enabled using the “enable mpls rsvp-te summary-refresh” command for a configured summary-refresh-time to actually be used.

Example

The following command sets the RSVP-TE hello time to 5 seconds on all MPLS interfaces:

```
configure mpls rsvp-te timers session hello-time 5 vlan all
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure mpls static lsp transport

```
configure mpls static lsp lsp_name transport [ip-traffic [allow | deny]
| vpn-traffic [allow {all | assigned-only} | deny]]
```

Description

Configures the type of traffic that can be transported across a static ingress LSP.

Syntax Description

<i>lsp_name</i>	Identifies the static LSP to be configured.
ip-traffic [allow deny]	Specifies whether IP traffic is to be allowed or denied access to the LSP.
vpn-traffic [allow {all assigned-only} deny]	Specifies whether VPN traffic is to be allowed or denied access to the LSP. The optional assigned-only keyword limits the transport of VPN traffic to only those VPLS instances that are explicitly configured to use the specified LSP.

Default

N/A.

Usage Guidelines

This command has no effect if the named LSP is a transit or egress LSP. By default, IP traffic and VPN traffic are set to deny for a newly created static LSP. The transport configuration options are independent. For example, if VPN traffic is set to allow and IP traffic is set to deny, then no routed IP traffic is transported across the LSP, but the LSP can still transport VPN traffic. When configured to deny for IP traffic, the specified LSP cannot be configured as an IP next hop for a default or static route.

Example

The following command configures a static LSP to transport IP traffic and all VPN traffic:

```
configure mpls static lsp lsp598 transport ip-traffic allow vpn-traffic allow all
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure mpls static lsp

```
configure mpls static lsp lsp_name [{egress [egress_label |
implicit-null] egress-vlan evlan_name next-hop ipaddress} {ingress
ingress_label {ingress-vlan ivlan_name}}]
```

Description

Configures the ingress and egress segments of a static LSP.

Syntax Description

<i>lsp_name</i>	Identifies the static LSP to be configured.
<i>egress_label</i>	Specifies the egress label for the LSP. The supported range is x7FC00 to x803FF. The egress label should match the corresponding ingress label of the next hop. There is no egress label at the egress LSR of a static LSP.
egress implicit-null	If PHP is supported, an LSR can be configured to use the implicit-null label for LSPs that terminate at the next-hop LER.
<i>evlan_name</i>	Specifies the egress VLAN for the LSP.
<i>ipaddress</i>	Specifies the IP address for the next-hop router along the static LSP.
<i>ingress_label</i>	Identifies the ingress label for this LSP. The supported range is x7FC00 to x7FFFF at transit LSRs and 0x80000 to 0x803FF at destination LSRs. The ingress label should match the corresponding egress label of the previous hop. There is no ingress label at the ingress LSR of a static LSP.
<i>ivlan_name</i>	When an ingress label is specified, this argument optionally specifies the ingress VLAN for the LSP.

Default

N/A.

Usage Guidelines

The ingress and egress segments can be configured any time before enabling the LSP. At the ingress LER, only the egress segment is configured and at the egress LER, only the ingress segment is configured. For LSPs that transit an LSR, it is mandatory to configure both ingress and egress segments. On any given LSR, the ingress label, if present, must match the egress label on the upstream LSR and the egress label must match the ingress label of the downstream LSR. Once configured, any change to the ingress or egress segments requires administratively disabling the LSP first. If the next-hop IP address is not within the subnet as defined by the interface VLAN name, the configuration is rejected.

Example

The following command configures a static LSP on an ingress LSR:

```
configure mpls static lsp lsp1 egress 0x7fc01 egress-vlan v50 next-hop 50.0.0.2
```

The following command configures a static LSP on a transit LSR:

```
configure mpls static lsp lsp1 egress 0x80001 egress-vlan v100 next-hop 100.0.0.2 ingress 0x7fc01 ingress-vlan v50
```

The following command configures a static LSP on an egress LSR:

```
configure mpls static lsp lsp1 ingress 0x80001 ingress-vlan v100
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure mrp ports timers

```
configure mrp ports [ port_list | all ] timers [{extended-refresh
[extended_refresh | off]} {join join_msec } {leave leave_msec }
{leave-all leave_all_msec } {periodic [periodic_msec | off]}]
```

Description

This command sets the join, leave, leave all, periodic, and extended-refresh timer values for a list of ports. The unit value is in milliseconds. The join timer, leave all timer, and periodic timer are started for each MRP application per port. The leave timer is started for each state machine that is in LV (leave) state. The default values for join, leave, leave-all, are 200, 600, and 10000, respectively. The default values for join, leave, leave-all, periodic and extended-refresh timers are 200, 600, 10000, 1000, and 0 milliseconds, respectively.

Syntax Description

mrp	Multiple Registration Protocol.
ports	Ports.
<i>port_list</i>	Port list separated by a comma or "-" type="portlist_t".
all	All ports.
timers	Multiple Registration Protocol timers.

extended-refresh	Timer value to use in place of regular leave timer, only in cases when leave-all is received or sent.
<i>extended_refresh_msec</i>	Extended refresh timer value in milliseconds (range is 600 ms to 300000 ms, default is 10000 ms).
join	The time interval to delay sending MRP advertisements.
<i>join_msec</i>	Join timer value in milliseconds (range is 0 ms to 500 ms, default is 200 ms).
leave	The time interval to wait in the leaving state before transitioning to the empty state.
<i>leave_msec</i>	Leave timer value in milliseconds (range is 600 ms to 3000 ms, default is 600 ms).
leave-all	The time interval used to control the frequency of "leave all" messages.
<i>leave_all_msec</i>	Leave All timer value in milliseconds (range is 5000 ms to 20000 ms, default is 10000 ms).
periodic	The time interval between two periodic events.
<i>periodic_msec</i>	Periodic timer value in milliseconds (range is 1000ms to 300000 ms, default is 1000 ms); type="uint32_t".
off	Turn off timer.
refresh	Timer value to use in place of regular timer, only in cases when leave-all is received or sent.
auto-refresh	Automatically calculate timer values based on number of talkers and listeners.
<i>refresh_msec</i>	Refresh timer value in milliseconds (range is 600ms to 300000ms, default is 0ms (off)).

Default

The default values for join, leave, leave-all, are 200, 600, and 10000, respectively. The default values for join, leave, leave-all, periodic and extended-refresh timers are 200, 600, 10000, 1000, and 0 milliseconds, respectively.

Usage Guidelines

This command is used to set the join, leave, and leave-all timer values for a list of ports. The unit value is in milliseconds. The join timer and leave all timer are started for each MRP application per port. The leave timer is started for each state machine that is in LV (leave) state. The default values for these timers are 200, 600, and 10000, respectively.

```
configure mrp ports 4 timers join 300
configure mrp ports all timers leave-all 15000
configure mrp ports all timers join 300 leave-all 15000
```

History

This command was first available in ExtremeXOS 15.3.

The extended-refresh and period timer options were added in 15.3.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure msdp as-display-format

```
configure msdp as-display-format [asdot | asplain]
```

Description

Configures the AS number format displayed in show commands.

Syntax Description

asdot	Specifies the ASDOT format.
asplain	Specifies the ASPLAIN format.

Default

N/A.

Usage Guidelines

The ASPLAIN and ASDOT formats are described in RFC 5396, Textual Representation of Autonomous System (AS) Numbers.

Example

The following command selects the ASDOT 4-byte AS number format:

```
configure msdp as-display-format asdot
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the *MSDP* feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure msdp max-rejected-cache

```
configure msdp max-rejected-cache max-cache {vr vrname}
```

Description

Configures the maximum limit on rejected SA cache entries that an *MSDP* router will store in its database.

Syntax Description

<i>max-cache</i>	Specifies the maximum number of rejected SA cache entries that the MSDP router will store in its database. To remove the limit, enter 0 (zero) for the <i>max-cache</i> value.
<i>vrname</i>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

By default, the maximum cache entries stored is zero. That is, rejected SA cache entries are not stored. Any SA cache entries that are stored and not refreshed for six minutes are removed.

Usage Guidelines

SA cache are rejected because of:

- Peer-RPF failure
- Policy denied

When a previously rejected SA cache entry is accepted because of an RP reachability change or policy rule change, the rejected SA cache entry is moved to the accepted SA cache list.

By default, rejected SA cache entries are discarded. You can configure a limit for rejected cache entries to store them, which will help debug/diagnose some issues; however, it consumes extra memory.

Example

The following command sets the maximum rejected cache limit to 100 for an MSDP router:

```
configure msdp max-rejected-cache 100
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure msdp originator-id

```
configure msdp originator-id ip-address {vr vrname}
```

Description

Configures the originator ID for an *MSDP* router. The originator ID is the RP address you want to use (instead of the default) in locally originated SA messages.

Syntax Description

<i>ip-address</i>	Specifies the RP address to use in locally originated SA messages. To unconfigure an originator ID (that is, to use the default RP address), enter the IP address 0.0.0.0.
<i>vrname</i>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

By default, the RP address is used as the originator ID in locally originated SA messages.

Usage Guidelines

Use this command to override the default RP address used in SA messages. Because only RPs and MSDP border routers originate SAs, there are times when it is necessary to change the ID used for this purpose. The originator ID address must be one of the interface addresses on the MSDP router.

You can configure the MSDP originator ID only when MSDP is disabled globally.

To remove an originator ID, enter the IP address 0.0.0.0.

Example

The following example configures the originator ID for an MSDP router:

```
configure msdp originator-id 10.203.134.1
```

The following example unconfigures the originator ID for an MSDP router:

```
configure msdp originator-id 0.0.0.0
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure msdp peer default-peer

```
configure msdp peer [remoteaddr | all] default-peer {default-peer-policy  
  filter-name} {vr vrname}
```

Description

This command configures a default or static RPF peer from which all *MSDP* SA messages are accepted. To remove the default peer, enter the `configure msdp peer no-default-peer` command.

Syntax Description

<i>filter-name</i>	Specifies the name of the policy filter associated with the default peer. The peer will be the default peer for all SA entries that are permitted by the policy filter. If an SA message is allowed by the policy filter, it will be accepted. Otherwise, the SA message has to go through the regular RPF-check. The static peer RPF check is the last step in peer RPF algorithm. So, if an SA message is denied by the default peer policy, ultimately the SA message will be rejected by MSDP.
peer all	Specifies all MSDP peers.
<i>remoteaddr</i>	Specifies the IP address of the MSDP peer.
<i>vrname</i>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

By default, no static RPF peer is configured.

The **default-peer-policy** keyword specifies the name of the policy filter associated with the default peer. You can configure multiple default peers with different policies. If no policy is specified, then the current peer is the default RPF peer for all SA messages.

Usage Guidelines

Configuring a default peer simplifies peer-RPF checking of SA messages. If the peer-RPF check fails, the default peer rule is applied to see if the SA messages should be accepted or rejected.

If a default peer policy is specified, the peer is the default peer only for the (Source, Group), or (S, G), that satisfies the policy. If the policy is not specified, then the default peer is used for all (S, G, RP).

You can configure multiple default peers on an MSDP router; however, all default peers must either have a default policy or not. A mix of default peers, with a policy and without a policy, is not allowed.

When configuring multiple default peer rules, follow these guidelines:

- When you enter multiple default-peer commands with the default-peer-policy keyword, you can use all the default peers at the same time for different RP prefixes.
- When you enter multiple default-peer commands without the default-peer-policy keyword, you can use a single active peer to accept all SA messages. If that peer goes down, then the next configured default peer accepts all SA messages. This configuration is typically used at a stub site.

You can use the following policy attributes in a default peer policy. All other attributes are ignored.

- Match:
 - multicast-group
 - multicast-source
 - pim-rp
- Set:
 - permit.
 - deny.

Example

The following example configures an MSDP peer with the IP address 192.168.45.43 as the default peer policy for "sales":

```
configure msdp peer 192.168.45.43 default-peer default-peer-policy sales
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure msdp peer description

```
configure msdp peer remoteaddr description {peer-description} {vr  
    vrname}
```

Description

Configures a name or description for an MSDP peer. This text is for display purposes only.

Syntax Description

<i>remoteaddr</i>	Specifies the IP address of the MSDP peer.
<i>peer-description</i>	Specifies the name or description of the MSDP peer. The maximum is 63 characters.
<i>vrname</i>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

By default, no name or description is specified.

Usage Guidelines

Use this command to configure a name or description to make an MSDP peer easier to identify. The description is visible in the output of the `show msdp peer` command.

To remove the description, use this command without a description string.

Example

The following example configures the name "internal_peer" to an MSDP peer:

```
configure msdp peer 192.168.45.43 description internal_peer
```

The following example removes the description from an MSDP peer:

```
configure msdp peer 192.168.45.43 description
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure msdp peer mesh-group

```
configure msdp peer [remoteaddr | all] mesh-group [mesh-group-name | none] [vr vrname]
```

Description

This command configures an *MSDP* peer to become a member of a mesh-group. To remove a peer from a mesh-group, enter the none CLI keyword for the mesh-group.

Syntax Description

<i>mesh-group-name</i>	Specifies the name of the MSDP mesh-group.
none	Removes a peer from a mesh-group.
peer all	Specifies all MSDP peers.
<i>remoteaddr</i>	Specifies the IP address of the MSDP peer.
<i>vrname</i>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

N/A.

Usage Guidelines

A mesh-group is a group of MSDP peers with fully meshed MSDP connectivity. Any SA messages received from a peer in a mesh-group are not forwarded to other peers in the same mesh-group.

Mesh-groups achieve two goals:

- Reduce SA message flooding.
- Simplify peer-RPF flooding.

Example

The following example configures an MSDP peer with the IP address 192.168.45.43 to become a member of a mesh-group called "intra":

```
configure msdp peer 192.168.45.43 mesh-group intra
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure msdp peer no-default-peer

```
configure msdp peer [remoteaddr | all] no-default-peer {vr vrname}
```

Description

This command removes a default peer.

Syntax Description

peer all	Specifies all <i>MSDP</i> peers.
<i>remoteaddr</i>	Specifies the IP address of the MSDP peer.
no-default-peer	Removes a default peer.
<i>vrname</i>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

N/A.

Usage Guidelines

None.

Example

The following command removes all MSDP peers:

```
configure msdp peer all no-default-peer
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure msdp peer password

```
configure msdp peer [remoteaddr | all] password {none | {encrypted}  
                  encrypted_tcp_password | tcp_password } {vr vrname}
```

Description

This command configures a TCP RSA Data Security, Inc. MD5 Message-Digest Algorithm password for an MSDP peer. This command enables TCP RSA Data Security, Inc. MD5 Message-Digest Algorithm authentication for a MSDP peer. When a password is configured, MSDP receives only authenticated MSDP messages from its peers. All MSDP messages that fail TCP RSA Data Security, Inc. MD5 Message-Digest Algorithm authentication are dropped.

Syntax Description

peer all	Specifies all MSDP peers.
<i>remoteaddr</i>	Specifies the IP address of the MSDP peer.
none	Removes the previously configured password.
encrypted	Encrypts the password for RSA Data Security, Inc. MD5 Message-Digest Algorithm authentication. To improve security, the password displays in encrypted format and cannot be seen as simple text. Additionally, the password is saved in encrypted format.
<i>tcpPassword</i>	Specifies the password to use for RSA Data Security, Inc. MD5 Message-Digest Algorithm authentication at the TCP level. The password must be an ASCII string with a maximum of 31 characters.
<i>vrname</i>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Defaults

By default, TCP RSA Data Security, Inc. MD5 Message-Digest Algorithm authentication is disabled for the MSDP peer.

Usage Guidelines

We recommend that you enable TCP RSA Data Security, Inc. MD5 Message-Digest Algorithm authentication for all MSDP peers to protect MSDP sessions from attacks. You can execute this command only when the MSDP peer is disabled or when MSDP is globally disabled on that VR.

Example

The following example configures a password for the MSDP peer with the IP address 192.168.45.43, which automatically enables TCP MD5 authentication:

```
configure msdp peer 192.168.45.43 password test123
```

The following command removes the password:

```
configure msdp peer 192.168.45.43 password none
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure msdp peer sa-filter

```
configure msdp peer [remoteaddr | all] sa-filter [in | out] [filter-name  
| none] {vr vr_name}
```

Description

This command configures an incoming or outgoing policy filter for SA messages.

Syntax Description

peer all	Specifies all <i>MSDP</i> peers.
<i>remoteaddr</i>	Specifies the IP address of the MSDP peer.
in	Associates the SA filter with inbound SA messages.
out	Associates the SA filter with outbound SA messages.

<i>filter-name</i>	Specifies the name of the policy associated with an SA filter. To remove an SA filter, enter the none CLI keyword instead of <i>filter-name</i> .
<i>vr_name</i>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

By default, no SA filter is configured for an MSDP peer. That is, incoming and outgoing SA messages are not filtered.

Usage Guidelines

This command configures an SA filter such that only a specified set of SA messages are accepted or sent to a peer. Note that an SA filter does not adversely impact the flow of SA request and response messages.

To remove an SA filter, enter the **none** CLI keyword instead of *filter-name*.

You can use the following policy attributes in an SA filter policy. All other attributes are ignored.

- Match:
 - multicast-group
 - multicast-source
 - pim-rp
- Set:
 - permit
 - deny

Example

The following example configures an incoming SA messages filter on an MSDP peer with the IP address 192.168.45.43:

```
configure msdp peer 192.168.45.43 sa-filter in allow_229
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure msdp peer sa-limit

```
configure msdp peer [remoteaddr | all] sa-limit max-sa {vr vr_name}
```

Description

This command allows you to limit the number of SA entries from an *MSDP* peer that the router will allow in the SA cache. To allow an unlimited number of SA entries, use 0 (zero) as the value for *max-sa*.

Syntax Description

peer all	Specifies all MSDP peers.
<i>remoteaddr</i>	Specifies the IP address of the MSDP peer.
<i>max-sa</i>	Specifies the maximum number of SA entries from an MSDP peer allowed in the SA cache. To specify an unlimited number of SA entries, use 0 (zero) as the value for <i>max-sa</i> .
<i>vr_name</i>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

By default, no SA entry limit is set. The router can receive an unlimited number of SA entries from an MSDP peer.

Usage Guidelines

You can use this command to prevent a distributed denial of service (DOS) attack. We recommend that you configure an MSDP SA limit on all MSDP peer sessions. Note that a rejected SA cache entry is not included in the number of SA cache entries received from a peer.

Example

The following example configures the SA entry limit of 500 for the MSDP peer with the IP address 192.168.45.43:

```
configure msdp peer 192.168.45.43 sa-limit 500
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure msdp peer source-interface

```
configure msdp peer [remoteaddr | all] source-interface [ipaddress |  
  any] {vr vrname}
```

Description

This command configures the source interface for the *MSDP* peer TCP connection.

Syntax Description

peer all	Specifies all MSDP peers.
<i>remoteaddr</i>	Specifies the IP address of the MSDP peer.
<i>ipaddress</i>	Specifies the IP address of the MSDP router interface to use on one end of a TCP connection. The <i>ipaddress</i> must be one of the MSDP router interface addresses; otherwise, the command fails and an error message displays.
any	Specifies to use any interface as one end of the TCP connection. The source interface is selected based on the IP route entry used to reach the MSDP peer. The egress interface that reaches the MSDP peer is used as the source interface for the TCP connection. Basically, this command removes the previously configured source interface of the MSDP peer.
<i>vrname</i>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Defaults

By default, the source interface is selected based on the IP route entry used to reach the MSDP peer. The egress interface that reaches the MSDP peer is used as the source interface for the TCP connection.

Usage Guidelines

You must first disable MSDP or the MSDP peer before using this command. We recommend that you configure a source interface for MSDP peers that are not directly connected. We also recommend using the loopback address as the MSDP peer connection endpoint.

Example

The following example configures a source interface for an MSDP peer with the IP address 192.168.45.43:

```
configure msdp peer 192.168.45.43 source-interface 60.0.0.5
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure msdp peer timer

```
configure msdp peer [remoteaddr | all] timer keep-alive keep-alive-sec  
hold-time hold-time-sec {vr vrname}
```

Description

The command configures the keep-alive and hold timer intervals of the *MSDP* peers.

Syntax Description

peer all	Specifies all MSDP peers.
<i>remoteaddr</i>	Specifies the IP address of the MSDP peer.
<i>keep-alive-sec</i>	Specifies the keep-alive timer interval in seconds. The range is 1–60 seconds.
<i>hold-time-sec</i>	Specifies the hold timer interval in seconds. The range is 3–75 seconds.
<i>vrname</i>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

By default, the:

- Keep-alive timer interval is 60 seconds.
- Hold timer interval is 75 seconds.
- SA timer interval is 60 seconds.

Usage Guidelines

You can use this command only when either MSDP or the MSDP peer is disabled. The hold timer interval must be greater than the keep-alive timer interval.

Example

The following example configures the keep-alive and hold timer intervals for the MSDP peer 55.0.0.83:

```
configure msdp peer 55.0.0.83 timer keep-alive 30 hold-time 60
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure msdp peer ttl-threshold

```
configure msdp peer [remoteaddr | all] ttl-threshold ttl {vr vrname}
```

Description

Configures the limit to which multicast data packets are sent in SA messages to an *MSDP* peer. If the time-to-live (TTL) in the IP header of an encapsulated data packet exceeds the TTL threshold configured, encapsulated data is not forwarded to MSDP peers.

Syntax Description

<i>remoteaddr</i>	Specifies the IP address of the MSDP peer on which to configure a TTL threshold.
all	Specifies all MSDP peers.
<i>ttl</i>	Specifies the TTL value. The range is 0-255. To restore the default value, enter a TTL value of 0 (zero).
<i>vrname</i>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

The default value is zero, meaning all multicast data packets are forwarded to the peer regardless of the TTL value in the IP header of the encapsulated data packet.

Usage Guidelines

This command allows you to configure a TTL value to limit multicast data traffic.

Example

The following example configures a TTL threshold of 5:

```
configure msdp peer 192.168.45.43 ttl-threshold 5
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure msdp sa-cache-server

```
configure msdp sa-cache-server remoteaddr {vr vr_name}
```

Description

Configures the *MSDP* router to send SA request messages to the MSDP peer when a new member becomes active in a group.

Syntax Description

<i>remoteaddr</i>	Specifies the IP address of the MSDP peer from which the local router requests SA messages when a new member becomes active in a group, and MSDP has no cache entry for the group in the local database.
<i>vr_name</i>	Specifies the name of the virtual router on which the MSDP cache server is configured. If a virtual router name is not specified, it is extracted from the current CLI context.

Default

By default, the router does not send SA request messages to its MSDP peers when a new member joins a group and wants to receive multicast traffic. The new member simply waits to receive SA messages, which eventually arrive.

Usage Guidelines

You can use this command to force a new member of a group to learn the current active multicast sources in a connected PIM-SM domain that are sending to a group. The router will send SA request messages to the specified MSDP peer when a new member joins a group and MSDP doesn't have a cache entry for that group in the local database. The peer replies with the information in an SA cache response message.



Note

An MSDP peer must exist before it can be configured as an SA cache server. The `configure msdp sa-cache-server` command accepts the value for *remoteaddr* only if it is an existing peer's IP address.

Example

The following example configures an MSDP cache server:

```
configure msdp sa-cache-server 172.19.34.5
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure msrp latency-max-frame-size

```
configure msrp [ latency-max-frame-size frame_size | [ ignore-latency-  
changes | talker-vlan-pruning ] [ on | off ] ]
```

Description

This command configures the system-wide MSRP variables.

Syntax Description

msrp	Multiple Stream Registration Protocol.
latency-max-frame-size	Maximum size of interfering frame (used in latency calculations).
<i>frame_size</i>	The maximum frame size in bytes (range 64 to 2000, default is 1522).
ignore-latency-changes	Ignore accumulated latency changes when evaluating first value change.
talker-vlan-pruning	Talker propagation is filtered on ports where <u>VLAN</u> does not exist.
on	Turn on.
off	Turn off.

Default

1522.

Usage Guidelines

Use this command to configure the system-wide MSRP variables.

Example

```
# configure msrp latency-max-frame-size 100
```

History

This command was first available in ExtremeXOS 15.3. The ignore-latency-changes, talker-vlan-pruning, and on | off options were added in 15.3.2.

Platform Availability

This command is available on platforms that support the AVB feature pack license and have it installed. For complete information about software licensing, including how to obtain and upgrade your license, and which platforms support the AVB feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure msrp ports sr-pvid

```
configure msrp ports [port_list | all] sr-pvid vlan_tag
```

Description

Specifies the default VLAN ID on the port for MSRP data stream. The sr-pvid serves as a recommendation to connected AVB devices; AVB devices may still use other VLAN IDs if they are configured to do so.

Syntax Description

msrp	Multiple Stream Registration Protocol
<i>port_list</i>	List of ports in the switch.
all	All the ports in the switch.
sr-pvid	Default VLAN Identifier for stream-related traffic.
<i>vlan_tag</i>	VLAN ID ranging from 1 to 4094 (default is 2).

Default

2.

Usage Guidelines

Use this command to specify the default VLAN ID on the port for MSRP data streams. The sr-pvid serves as a recommendation to connected AVB devices; AVB devices may still use other VLAN IDs if they are configured to do so.

Example

```
# configure msrp ports 1,2,3 sr-pvid 2
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the AVB feature pack license and have it installed. For complete information about software licensing, including how to obtain and upgrade your license,

and which platforms support the AVB feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure msrp ports traffic-class delta-bandwidth

```
configure msrp ports [port_list | all] traffic-class [A | B] delta-  
bandwidth percentage
```

Description

Configures delta-bandwidth value per traffic class per MSRP port.

Syntax Description

msrp	Multiple Stream Registration Protocol.
<i>port_list</i>	List of ports in the switch.
traffic-class	Traffic class.
A	Traffic class A.
B	Traffic class B.
delta-bandwidth	Delta-bandwidth percentage (range 0 to 100, default 75 for class A, 0 for class B).

Default

Class A: 75, Class B: 0.

Usage Guidelines

The delta bandwidth configuration limits the amount of bandwidth that can be used by the given stream reservation class. Each class is allowed to use a maximum of its delta bandwidth plus the delta bandwidth configured for each of the higher classes. For example, if the delta bandwidth for classes A and B are configured to 10 and 10 respectively, class A streams can use up to 10 percent of the link bandwidth, and class B streams can use up to 20 percent of the link bandwidth. The sum of the class A and B delta bandwidth values must be less than 100 percent.

Example

```
# configure msrp ports all traffic-class A delta-bandwidth 50  
# configure msrp ports 1-5 traffic-class B delta-bandwidth 0
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the AVB feature pack license and have it installed. For complete information about software licensing, including how to obtain and upgrade your license, and which platforms support the AVB feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure msrp sharing

```
configure msrp sharing [all | port_list] bandwidth [cumulative | single-port] percentage
```

Description

This command configures the [LAG](#) bandwidth mode as either cumulative or master-port only.

Syntax Description

all	All the ports in the switch.
<i>port_list</i>	Port list separated by a comma or -
cumulative	Use bandwidth of a single port, plus a percentage of bandwidth of every other LAG member port in the group.
single-port	Use bandwidth of a single port only for the entire LAG.
<i>percentage</i>	Percentage of bandwidth of each LAG port to be added to master port bandwidth.

Default

Single-port.

Usage Guidelines

If cumulative mode is selected, the percentage is also configured.

Example

This CLI command displays bandwidth information of an MSRP port.

```
# show msrp ports bandwidth
Port      Port Class  Delta      Maximum    Reserved   Available
      Speed
-----
5ab        0 M A      75.00%    0.00%      0.00%      0.00%
           B      0.00%    0.00%      0.00%      0.00%
*21ab     1000 M A    75.00%    75.00%      0.00%      75.00%
           B      0.00%    75.00%      0.00%      75.00%
Flags: (*) Active, (!) Administratively disabled,
      (a) SR Class A allowed, (b) SR Class B allowed.
```

History

This command was first available in ExtremeXOS 21.1.

Platform Availability

This command is available on platforms that support the AVB feature pack license and have it installed. For complete information about software licensing, including how to obtain and upgrade your license, and which platforms support the AVB feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure msrp timers first-value-change-recovery

```
configure msrp timers first-value-change-recovery
    [first_value_change_msec | off]
```

Description

This command configures MSRP first value change recovery timer, or disables the timer. If configured, the system waits until the configured timer value before allowing recovery of streams from first value change failure. If disabled, the system does not recover from first value change failure.

Syntax Description

msrp	Multiple Stream Registration Protocol.
timers	Multiple Stream Registration Protocol timers.
first-value-change-recovery	The time interval to wait to allow recovery of stream from first value change failure.
first_value_change_msec	First Value Change Recovery time in milliseconds (range is 10000 ms to 5400000 ms, default is 30000 ms); type="uint32_t"; range="[10000, 5400000]".
off	Turn off first value change recovery timer, and do not recover from first value change failure.

Default

30000 ms.

Usage Guidelines

Use this command to allow streams to recover from first value change failure.

Example

```
# configure msrp timers first-value-change recovery 20000
# configure msrp timers first-value-change recovery off
```

History

This command was first available in ExtremeXOS 15.3.2.

Platform Availability

This command is available on platforms that support the AVB feature pack license and have it installed. For complete information about software licensing, including how to obtain and upgrade your license, and which platforms support the AVB feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure mstp format

```
configure mstp format format_identifier
```

Description

Configures the number used to identify the *MSTP* BPDUs sent in the MSTP region.

Syntax Description

<i>format_identifier</i>	Specifies a number that MSTP uses to identify all BPDUs sent in the MSTP region. The default is 0. The range is 0 to 255.
--------------------------	---

Default

The default value used to identify the MSTP BPDU is 0.

Usage Guidelines

For a switch to be part of an MSTP region, you must configure each switch in the region with the same MSTP configuration attributes, also known as MSTP region identifiers. These identifiers consist of the following:

- Region Name—The name of the MSTP region.
- Format Selector—The number used to identify the format of MSTP BPDUs. The default is 0.
- Revision Level—This identifier is reserved for future use; however, the switch uses and displays a default of 3.

You can configure only one MSTP region on the switch at any given time.

The switches contained in a region transmit and receive BPDUs that contain information relevant to only that MSTP region. By having devices look at the region identifiers, MSTP discovers the logical boundary of a region.

If you have an active MSTP region, Extreme Networks recommends that you disable all active STPDs in the region before modifying the value used to identify MSTP BPDUs on all participating switches.

Example

The following command configures the number 2 to identify the MSTP BPDUs sent within an MSTP region:

```
configure mstp format 2
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure mstp region

```
configure mstp region regionName
```

Description

Configures the name of an *MSTP* region on the switch.

Syntax Description

<i>regionName</i>	Specifies a user-defined name for the MSTP region. May be up to 32 characters.
-------------------	--

Default

By default, the switch uses the MAC address of the switch to generate an MSTP region.

Before you configure the MSTP region, it also has the following additional defaults:

- MSTP format Identifier—0.
- MSTP Revision Level—3.

Usage Guidelines

The maximum length for a name is 32 characters. Names can contain alphanumeric characters and underscores (_) but cannot be any reserved keywords, for example, mstp. Names must start with an alphabetical character, for example, a, Z.

By default, the switch uses the unique MAC address of the switch to generate an MSTP region. Since each MAC address is unique, every switch is in its own region by default.

For multiple switches to be part of an MSTP region, you must configure each switch in the region with the same MSTP configuration attributes, also known as MSTP region identifiers. These identifiers consist of the following:

- Region Name—The name of the MSTP region.
- Format Selector—The number used to identify the format of MSTP BPDUs. The default is 0.
- Revision Level—This identifier is reserved for future use; however, the switch uses and displays a default of 3.

You can configure only one MSTP region on the switch at any given time.

The switches inside a region exchange BPDUs that contain information for MSTIs. The switches connected outside of the region exchange CIST information. By having devices look at the region identifiers, MSTP discovers the logical boundary of a region.

If you have an active MSTP region, we recommend that you disable all active STPDs in the region before renaming the region on all of the participating switches.

Viewing MSTP Information

To view the MSTP configuration on the switch, use the `show stpd` command. Output from this command contains global MSTP settings, including the name of the MSTP region, the number or tag that identifies all of the BPDUs sent in the MSTP region, and the reserved MSTP revision level. If configured, the output also displays the name of the Common and Internal Spanning Tree (CIST), and the number of Multiple Spanning Tree Instances (MSTIs).

Example

The following example creates an MSTP region named purple:

```
configure mstp region purple
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure mstp revision

```
configure mstp revision revision
```

Description

Configures the revision number of the *MSTP* region.

Syntax Description

<i>revision</i>	This parameter is reserved for future use.
-----------------	--

Default

The default value of the revision level is 3.

Usage Guidelines

Although this command is displayed in the CLI, it is reserved for future use. Please do not use this command.

If you accidentally configure this command, remember that each switch in the region must have the same MSTP configuration attributes, also known as MSTP region identifiers. These identifiers consist of the following:

- Region Name—The name of the MSTP region.
- Format Selector—The number used to identify the format of MSTP BPDUs. The default is 0.
- Revision Level—An unsigned integer encoded within a fixed field of 2 octets that identifies the revision of the current MST configuration. MSTP revision level can be set from 0 to 65536, with the default being 3. The revision number is not incremented automatically each time that the MST configuration is committed.

Example

The following command returns the MSTP revision number to 3, the default revision number:

```
configure mstp revision 3
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure mvr add receiver

```
configure mvr vlan vlan-name add receiver port port-list
```

Description

Configures a port to receive MVR multicast streams.

Syntax Description

<i>vlan-name</i>	Specifies a VLAN name.
<i>port-list</i>	A list of ports or slots and ports.

Default

N/A.

Usage Guidelines

This command is used to add a group of virtual ports for multicast forwarding through MVR. By default, some ports on non-MVR VLANs (router ports, primary and secondary [EAPS](#) ports), are excluded from the MVR cache egress list. This command is used to override these rules, so that if valid [IGMP](#) memberships are received, or a router is detected, streams are forwarded out on the ports.

Example

The following example adds the ports 1:1 and 1:2 of VLAN v1 to MVR for forwarding:

```
configure mvr vlan v1 add receiver port 1:1-1:2
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MVR feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure mvr add vlan

```
configure mvr add vlan vlan-name
```

Description

Configures a [VLAN](#) as an MVR VLAN.

Syntax Description

<i>vlan-name</i>	Specifies a VLAN name.
------------------	------------------------

Default

N/A.

Usage Guidelines

Configures MVR on the specified VLAN. When a multicast stream in the specified MVR address range is received on the VLAN, it is leaked to all other VLAN ports where the corresponding *IGMP* join message is received. By default, the entire multicast address range 224.0.0.0/4, except for the multicast control range 224.0.0.0/24 is used for MVR. To change the MVR address range, use the following command:

```
configure mvr vlan vlan-name mvr-address {policy-name | none}
```

Example

The following example configures VLAN v1 as an MVR VLAN:

```
configure mvr add vlan v1
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MVR feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure mvr delete receiver

```
configure mvr vlan vlan-name delete receiver port port-list
```

Description

Configures a port not to receive MVR multicast streams.

Syntax Description

<i>vlan-name</i>	Specifies a <i>VLAN</i> name.
<i>port-list</i>	A list of ports or slots and ports.

Default

N/A.

Usage Guidelines

This command is used to delete a group of virtual ports for multicast forwarding through MVR. After using this command, the ports revert to the default forwarding rules.

Example

The following example deletes the ports 1:1 and 1:2 of VLAN v1 to MVR for forwarding:

```
configure mvr vlan v1 delete receiver port 1:1-1:2
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MVR feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure mvr delete vlan

```
configure mvr delete vlan vlan-name
```

Description

Deletes a VLAN from MVR.

Syntax Description

<i>vlan-name</i>	Specifies a VLAN name.
------------------	------------------------

Default

N/A.

Usage Guidelines

Removes MVR from the specified VLAN.

Example

The following example configures VLAN v1 as a non-MVR VLAN:

```
configure mvr delete vlan v1
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MVR feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure mvr mvr-address

```
configure mvr vlan vlan-name mvr-address {policy-name | none}
```

Description

Configures the MVR address range on a VLAN.

Syntax Description

<i>vlan-name</i>	Specifies a VLAN name.
<i>policy-name</i>	Specifies a policy file.

Default

The default address range is 224.0.0.0/4 (all multicast addresses), but excluding 224.0.0.0/24 (the multicast control range).

Usage Guidelines

If no policy file is specified (the **none** option), the entire multicast address range 224.0.0.0/4, except for the multicast control range 224.0.0.0/24 is used for MVR.

MVR must first be configured on the VLAN before using this command.

If the policy is later refreshed, groups denied and newly allowed groups in the policy are flushed from fast path forwarding. This allows synching existing channels with the new policy, without disturbing existing channels.

The following is a sample policy file mvrpol.pol. This policy configures 236.1.1.0/24 as the MVR address range. Any address outside this range has the standard switching behavior on an MVR VLAN.

```
Entry extremel {  
    if match any {  
        nlri 236.1.1.0/24 ;  
    }  
    then {  
        permit ;  
    }  
}
```

Example

The following example configures the MVR address range specified in the policy file mvrpol.pol for the VLAN v1:

```
configure mvr vlan v1 mvr-address mvrpol
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MVR feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure mvr static group

```
configure mvr vlan vlan-name static group {policy-name | none}
```

Description

Configures the MVR static group address range on a VLAN.

Syntax Description

<i>vlan-name</i>	Specifies a VLAN name.
<i>policy-name</i>	Specifies a policy file.

Default

By default, all the MVR group addresses work in static mode.

Usage Guidelines

If no policy file is specified (the **none** option), the entire multicast address range 224.0.0.0/4, except for the multicast control range 224.0.0.0/24, is used for static groups in MVR.

MVR must first be configured on the VLAN before using this command.

The following is a sample policy file mvrpol.pol. This policy configures 236.1.1.0/24 as the MVR static group address range. Any MVR addresses outside this range are dynamically registered through IGMP. An MVR VLAN will proxy join only for addresses that are not in the static group. If you want all the multicast groups to be dynamic, use a policy file with this command that denies all multicast addresses.

```
Entry extremel {  
    if match any {  
        nlri 236.1.1.0/24 ;  
    }  
    then {
```

```
        permit ;
    }
}
```

Example

The following example configures the MVR static group address range specified in the policy file mvrpol.pol for the VLAN v1:

```
configure mvr vlan v1 static group mvrpol
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MVR feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure mvrp stpd

```
configure mvrp stpd stpd_name
```

Description

Configures the STP domain to use for dynamically created VLANs.

Syntax Description

mvrp	Multiple VLAN Registration Protocol.
stpd	The STP domain used for MVRP.
<i>stpd_name</i>	The STP domain the VLAN is to be associated. All ports of the domain will be advertised, when this VLAN gets registered.

Default

s0.

Usage Guidelines

Use this command to configure the STP domain used for MVRP.

Example

The following example configures the default STP domain for MVRP to "stpd2":

```
configure mvrp stpd stpd2
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure mvrp tag vlan_tag ports [port_list |all] registration [forbidden |normal]

```
configure mvrp tag vlan_tag ports [port_list |all] registration
[forbidden |normal ]
```

Description

This command is used for per port setting for the [VLAN](#) registration. If the global registration is forbidden, ports cannot be added to any VLAN dynamically irrespective of the per-port setting. So for ports to be registered, the global and the per-port setting both should be "normal", which is the default value.

Syntax Description

mvrp	Multiple VLAN Registration Protocol.
tag	The 802.1Q VLAN ID.
<i>vlan_tag</i>	VLAN ID ranging from 1 to 4094; type=uint16_t"; range="[1,4094]".
ports	Ports.
<i>port_list</i>	Port list separated by a comma or "-"; type="portlist_t";
all	All ports.
registration	Whether port can be added dynamically to the VLAN.
forbidden	Port cannot be added dynamically to the VLAN.
normal	Port can be added dynamically to the VLAN.

Default

Normal.

Usage Guidelines

Use this command to control dynamic addition of ports to VLANs.

Example

```
configure mvrp tag 2 ports 2,3,4 registration forbidden
configure mvrp tag 2 ports all registration normal
```

History

This command was first available in ExtremeXOS 15.3.

The registration option, and forbidden and normal keywords were added in 15.3.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure mvrp tag ports transmit

```
configure mvrp tag vlan_tag ports [port_list | all] transmit [on | off ]
```

Description

Controls whether the given VLAN ID may be advertised in MVRP messages transmitted on the given set of ports.

Syntax Description

mvrp	Multiple VLAN Registration Protocol.
tag	The 802.1Q VLAN ID.
transmit	When enabled, MVRP message are sent on the ports.
on	Transmission of MVRP messages are enabled on the port(s) for the given tag.
off	Transmission of the MVRP messages are disabled on the port(s) for the given tag.

Default

Transmit on.

Usage Guidelines

Use this command to control whether the given VLAN ID may be advertised in MVRP messages transmitted on the given set of ports.

Example

The following command configures transmit off for VLAN ID 100 on all MVRP ports:

```
configure mvrp tag 100 ports all transmit off
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure mvrp vlan auto-creation

```
configure mvrp vlan auto-creation [on | off]
```

Description

Enables or disables the dynamic VLAN creation feature of MVRP.

Syntax Description

mvrp	Multiple VLAN Registration Protocol.
auto-creation	When enabled, results in VLANs added dynamically on the switch through MVRP.
on	Enable auto-creation.
off	Disable auto-creation.

Default

Enabled.

Usage Guidelines

Use this command to enable or disable the dynamic VLAN creation of MVRP. By default, auto-creation is enabled. If disabled, the switch may participate in the MVRP protocol, and advertised static VLANs, but will not dynamically create VLANs.

Example

The following command enables MVRP VLAN auto creation:

```
configure mvrp vlan auto-creation on
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure mvrp vlan registration

```
configure mvrp vlan registration forbidden | normal
```

Description

This command is a global system setting. If global registration is forbidden, ports cannot be added to any VLAN dynamically.

Syntax Description

mvrp	Multiple VLAN Registration Protocol.
vlan	VLAN.
registration	Whether all ports can be added to new dynamic VLANs. This can be overridden by static port addition to VLAN.
forbidden	Ports cannot be added dynamically to the VLAN. This can be overridden by static port addition.
normal	Ports can be added dynamically to the VLAN (default).

Default

Normal.

Usage Guidelines

Use this command to set global registration. If global registration is forbidden, ports cannot be added to any VLAN dynamically.

Example

The following command allows ports to be added dynamically to the VLAN:

```
configure mvrp vlan registration normal
```

History

This command was first available in ExtremeXOS 15.3.

The **registration** keyword was first available in ExtremeXOS 15.3.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure neighbor-discovery cache add

```
configure neighbor-discovery cache {vr vr_name} add [ipv6address |  
  scoped_link_local] mac
```

Description

Adds a static entry to the neighbor cache.

Syntax Description

<i>vr_name</i>	Specifies a VR or VRF.
<i>ipv6address</i>	Specifies an IPv6 address.
<i>scoped_link_local</i>	Specifies a scoped, link-local address.
<i>mac</i>	Specifies a MAC address.

Default

If you do not specify a VR or VRF, the current VR context is used.

Usage Guidelines

This command adds static entries to the neighbor cache.

Example

The following example adds a static entry to the neighbor cache:

```
configure neighbor-discovery cache add fe80::2315%default 00:11:22:33:44:55
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure neighbor-discovery cache delete

```
configure neighbor-discovery cache {vr vr_name} delete [ipv6address |  
  scoped_link_local]
```

Description

Deletes a static entry from the neighbor cache.

Syntax Description

<i>vr_name</i>	Specifies a VR or VRF.
<i>ipv6address</i>	Specifies an IPv6 address.
<i>scoped_link_local</i>	Specifies a scoped, link-local address.

Default

If you do not specify a VR or VRF, the current VR context is used.

Usage Guidelines

This command deletes static entries from the neighbor cache.

Example

The following example deletes a static entry from the neighbor cache:

```
configure neighbor-discovery cache delete fe80::2315%default
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure neighbor-discovery cache locktime

```
configure neighbor-discovery cache {vr vr_name}{locktime locktime}
```

Description

Sets the time before a new entry can replace an old entry in the Neighbor Discovery Protocol (NDP) cache of neighbor IPv6 addresses\MAC addresses.

Syntax Description

vr	Specifies setting a VR or VRF.
<i>vr_name</i>	Specifies the name of the VR or VRF.

locktime	Specifies setting a time before a new entry can replace an old entry.
<i>locktime</i>	Sets the locktime value in milliseconds with a range of 0–30,000. Default is 1,000 milliseconds.

Default

The default locktime is 1,000 milliseconds.

Example

The following example sets the locktime to 5,000 milliseconds:

```
configure neighbor-discovery cache locktime 5000
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure neighbor-discovery cache max_entries

```
configure neighbor-discovery cache max_entries max_entries
```

Description

Configures the maximum allowed IPv6 neighbor entries.

Syntax Description

<i>max_entries</i>	Specifies the maximum allowed IPv6 neighbor entries. The range is 1 to 49,152.
--------------------	--

Default

8,192.

Usage Guidelines

For ExtremeXOS 30.1, the maximum configurable limit for neighbor discovery maximum entries is changed to 49,152 for all platforms. A message appears if the configured value exceeds the theoretical hardware maximum limit depending on the platform.

Example

The following example sets the maximum allowed IPv6 neighbor entries to 512:

```
configure neighbor-discovery cache max_entries 512
```

History

This command was first available in ExtremeXOS 12.4.

Per virtual router capability was deprecated and the maximum configurable limit set to 49,152 in ExtremeXOS 30.1.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure neighbor-discovery cache max_pending_entries

```
configure neighbor-discovery cache max_pending_entries
    max_pending_entries
```

Description

Configures the maximum number of pending IPv6 neighbor entries.

Syntax Description

<i>max_pending_entries</i>	Specifies the maximum number of pending IPv6 neighbor entries. The range is 1 to 4096.
----------------------------	--

Default

1,024.

Usage Guidelines

None.

Example

The following example sets the maximum number of pending IPv6 neighbor entries to 2,056:

```
configure neighbor-discovery cache max_pending_entries 2056
```

History

This command was first available in ExtremeXOS 12.4.

Per virtual router capability was deprecated in ExtremeXOS 30.1.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure neighbor-discovery cache reachable-time

```
configure neighbor-discovery cache {reachable-time [auto |  
    {reachable_time [seconds | milliseconds]}}
```

Description

Sets the value for Neighbor Discovery Protocol (NDP) reachable time

Syntax Description

reachable-time	Specifies setting the NDP reachable time.
<i>reachable_time</i>	Sets the value for the NDP reachable time (range is 1-1,474,515,000 millisecond or 1-1,474,515 second).
auto	Specifies having the NDP reachable time set automatically to 3/4 of the configured NDP timeout (default).
milliseconds	When setting the reachable time value, specifies milliseconds as the time unit (range is 1-1,474,515,000).
seconds	When setting the reachable time value, specifies seconds (range is 1-1,474,515) as the time unit (default).

Default

The default setting is for the reachable time to be set automatically to 3/4 of the configured NDP timeout. If you set the time manually, the default unit of measure for the value is seconds.

Example

The following example sets the reachable time to 500,000 seconds:

```
configure neighbor-discovery cache reachable-time 500000 seconds
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure neighbor-discovery cache retransmit-time

```
configure neighbor-discovery cache {retransmit-time retransmit_time}
```

Description

Sets the value for Neighbor Discovery Protocol (NDP) retransmit time

Syntax Description

retransmit-time	Specifies setting the retransmit time.
<i>retransmit_time</i>	Sets the retransmit time value (range is 1–4,294,967 seconds or 1–4,294,967,295 milliseconds). The default is 1 second.
milliseconds	When setting the retransmit time value, specifies milliseconds as the time unit (range is 1–4,294,967,295).
seconds	When setting the retransmit time value, specifies seconds (range is 1–4,294,967) as the time unit (default).

Default

The default setting for the retransmit time is 1 second. The default unit of measure is seconds.

Example

The following example sets the retransmit time to 500,000 seconds:

```
configure neighbor-discovery cache retransmit-time 500000 seconds
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure neighbor-discovery cache timeout

```
configure neighbor-discovery cache {vr vr_name} timeout timeout
```

Description

Configures a timeout value for entries in the neighbor cache.

Syntax Description

<i>vr_name</i>	Specifies a VR or VRF.
<i>timeout</i>	Specifies a timeout value for neighbor cache entries. The range is 1 to 32767 minutes.

Default

20 minutes.

Usage Guidelines

None.

Example

The following example configures the neighbor cache timeout for 30 minutes:

```
configure neighbor-discovery cache timeout 30
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document./>

configure netlogin add mac-list

```
configure netlogin add mac-list [mac {mask} | default] {encrypted  
    {encrypted_password | password} {ports port_list}
```

Description

Adds an entry to the MAC address list for MAC-based network login.

Syntax Description

<i>mac</i>	Specifies the MAC address to add.
<i>mask</i>	Specifies the number of bits to use for the mask.
default	Specifies the default entry.
encrypted	Used to display encrypted form of password in configuration files. Do not use.
<i>password</i>	Specifies the password to send for authentication.
ports	Specifies the port or port list to use for authentication.

Default

If no password is specified, the MAC address will be used.

Usage Guidelines

Use this command to add an entry to the MAC address list used for MAC-based network login.

If no match is found in the table of MAC entries, and a default entry exists, the default will be used to authenticate the client. All entries in the list are automatically sorted in longest prefix order.

configure netlogin add mac-list default configuration is added by default when **enable netlogin mac** is configured.

Associating a MAC Address to a Port

You can configure the switch to accept and authenticate a client with a specific MAC address. Only MAC addresses that have a match for the specific ports are sent for authentication. For example, if you associate a MAC address with one or more ports, only authentication requests for that MAC addresses received on the port(s) are sent to the *RADIUS (Remote Authentication Dial In User Service)* server. The port(s) block all other authentication requests that do not have a matching entry. This is also known as secure MAC.

To associate a MAC address with one or more ports, specify the ports option when using the `configure netlogin add mac-list [mac {mask} | default] {encrypted} {password} {portsport_list}` command.

You must enable MAC-based network login on the switch and the specified ports before using this command. If MAC-based network login is not enabled on the specified port(s), the switch displays a warning message similar to the following:

```
WARNING: Not all specified ports have MAC-Based NetLogin enabled.
```

If this occurs, make sure to enable MAC-based network login.

Example

The following command adds the MAC address 10:20:30:40:50:60 with the password foo to the list:

```
configure netlogin add mac-list 10:20:30:40:50:60 password foo
```

The following command associates MAC address 10:20:30:40:50:70 with ports 2:2 and 2:3. This means authentication requests from MAC address 10:20:30:40:50:70 are only accepted on ports 2:2 and 2:3:

```
configure netlogin add mac-list mac 10:20:30:40:50:70 ports 2:2-2:3
```

History

This command was first available in ExtremeXOS 11.1.

The ports option was added in ExtremeXOS 11.3.

Default configuration when **enable netlogin mac** is entered was added in ExtremeXOS 31.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches..

configure netlogin add proxy-port

```
configure netlogin add proxy-port tcp_port {http | https}
```

Description

Configure the ports that will be hijacked and redirected for HTTP or HTTPS traffic.

Syntax Description

<i>tcp_port</i>	Specifies the port to be hijacked.
-----------------	------------------------------------

Default

HTTP traffic.

Usage Guidelines

This command allows you to configure the ports that will be hijacked and redirected for HTTP or HTTPS traffic. For each hijacked proxy port, you must specify whether the port is to be used for HTTP or HTTPS traffic.

No more than 5 such ports are supported in addition to ports 80 and ports 443. Attempts to add more than 5 ports generate an error.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure netlogin agingtime

```
configure netlogin agingtime minutes
```

Description

Lets you configure network login aging.

Syntax Description

<i>minutes</i>	Specifies the aging time in minutes.
----------------	--------------------------------------

Default

The default value is 5.

Usage Guidelines

Use this command to configure the aging time for network login. The aging time is the time after which learned clients that failed authentication or did not attempt to authenticate are removed from the system. This prevents the switch from keeping all clients ever seen on a network-login-enabled port.

The range can be from 0 to 3000, where 0 indicates no age out.

Example

The following command specifies an aging time of 15 minutes:

```
configure netlogin agingtime 15
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure netlogin allowed-refresh-failures

```
configure netlogin allowed-refresh-failures num_failures
```

Description

Sets the number refresh failures.

Syntax Description

<i>num_failures</i>	Specifies the number of refresh failures. The range is from 0 to 5.
---------------------	---

Default

The default is 0.

Usage Guidelines

This command allows you to set the number of refresh failures allowed. You can set the number of failures to be from between 0 to 5. The default value is 0.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure netlogin authentication database-order

```
configure netlogin [mac | web-based] authentication database-order
    [[radius] | [local] | [radius local] | [local radius]]
```

Description

Configures the order of database authentication protocols to use.

Syntax Description

mac	Specifies MAC-based authentication.
web-based	Specifies Web-based authentication.
radius	Specifies an authentication order from only the <i>RADIUS</i> database.
local	Specifies an authentication order from only the local database.
radius local	Specifies an authentication order of RADIUS database first, followed by local.
local radius	Specifies an authentication order of local database first, followed by RADIUS.

Default

By default, the authentication order is RADIUS, local-user database.

Usage Guidelines

As of ExtremeXOS 16.1, the functionality of this command is more consistent with management authentications. If RADIUS responds with a reject, then that reject is honored. The only time the local database is checked is when the RADIUS server does not respond.

Example

The following command sets the database authentication order to local-user database, RADIUS:

```
configure netlogin mac authentication database-order local radius
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure netlogin authentication failure vlan

```
configure netlogin authentication failure vlan vlan_name {ports
port_list}
```

Description

Configures authentication failure VLAN on network login enabled ports.

Syntax Description

<i>vlan_name</i>	Specifies the name of the authentication failure VLAN.
<i>port_list</i>	Specifies one or more ports or slots and ports. If the ports keyword is not used, the command applies to all ports.

Default

By default, authentication failure VLAN is configured on all network login enabled ports if no port is specifically configured.

Usage Guidelines

Use this command to configure authentication failure VLAN on network login enabled ports. When a supplicant fails authentication, it is moved to the authentication failure VLAN and is given limited access until it passes the authentication either through RADIUS or local. Depending on the authentication database order for that particular network login method (MAC, web or dot1x), the other database is used to authenticate the client. If the final result is an authentication failure and if the authentication failure VLAN is configured and enabled on that port, the client is moved to that location.

There four different authentication orders which can be configured per authentication method currently. They are:

- RADIUS.
- local.
- RADIUS, local.
- local, RADIUS.

In each case, you must consider the end result in deciding whether to authenticate the client in authentication failure VLAN or authentication service unavailable VLAN (if configured).

For example, when netlogin mac authentication database order is local, radius, if the authentication of a MAC client fails through a local database, RADIUS is used for authentication. If RADIUS also fails authentication, the client is moved to authentication failure VLAN. The same is true for all authentication database orders (radius,local; local,radius; radius; local).

If authentication through local fails, but passes through RADIUS, the client is moved to the appropriate destination VLAN.

If the local authentication fails and the RADIUS server is not available, the client is not moved to authentication failure VLAN.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure netlogin authentication protocol-order

```
configure netlogin authentication protocol-order [[dot1x [web-based |  
mac | cep]] | [mac [dot1x | web-based | cep]] | [web-based [dot1x |  
mac | cep]] | [cep [dot1x | web-based | mac]]]
```

Description

Globally configures the order of the Network Login (NetLogin) port's authentication protocols.

Syntax Description

dot1x	Configures the 802.1x authentication protocol preference.
mac	Configures the MAC-based authentication protocol preference.
web-based	Configures the web-based authentication protocol preference.
cep	Configure Convergence End Point (CEP) authentication protocol preference. CEP only appears as an option if policy is enabled.

Default

By default, the protocol precedence order for a NetLogin-enabled port is:

- Dot1x
- Web-based
- MAC
- CEP

Usage Guidelines

Web-based authentication occurs only when the port belongs to the NetLogin VLAN.

When you change the protocol precedence, the action for the current highest precedence protocol takes effect immediately if the client is authenticated by this protocol.

When you disable the highest precedence protocol on a port, the action for the next precedence protocol takes effect immediately if client is authenticated by this protocol.

CEP only appears as an option in the command if policy is enabled.

Example

The following example sets the protocol precedence order to Dot1x, Web-based, and MAC.

```
configure netlogin authentication protocol-order dot1x web-based mac cep
```

History

This command was first available in ExtremeXOS 15.7.1.

CEP option was added in ExtremeXOS 32.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure netlogin authentication service-unavailable vlan

```
configure netlogin authentication service-unavailable [{add} | {delete}  
| [{vlan vlan_name} {ports port_list {tagged | untagged}}}]
```

Description

Configures authentication service-unavailable VLAN on NetLogin-enabled ports.

Syntax Description

<i>vlan_name</i>	Specifies the name of the service-unavailable VLAN.
<i>port_list</i>	Specifies one or more ports or slots and ports. If the ports keyword is not used, the command applies to all ports.
add	Add service-unavailable VLAN to ports (default).
tagged	Configure port as tagged to the service-unavailable VLAN.
untagged	Configure port as untagged to the service-unavailable VLAN (default).
delete	Delete existing service-unavailable VLAN from ports.

Default

If a port is not specified, all NetLogin-enabled ports are applied.

If not specified, the command adds service-unavailable VLAN to ports by default.

If not specified, the ports are configured as untagged to the service-unavailable VLAN by default.

Usage Guidelines

This command configures authentication service-unavailable VLAN(s) on the specified NetLogin-enabled ports. Authentication service-unavailable VLAN is configured on all the NetLogin-enabled ports, if no port is specifically selected. When an authentication service is not available to authenticate the NetLogin clients, they are moved to the authentication service-unavailable VLAN(s) and are given limited access until the authentication service is available through [RADIUS](#).

Starting with ExtremeXOS 30.2, you can specify up to 10 service-unavailable VLANs per port.

As of ExtremeXOS 16.1, the functionality of this command is more consistent with management authentications. If RADIUS responds with a reject, then that reject is honored.

There are four different authentication orders that can be configured per authentication method currently. They are:

- RADIUS
- Local
- RADIUS, local
- Local, RADIUS

The service unavailable VLAN is used only when authentication order is "RADIUS". The authentication failure VLAN is used for all other modes (local; RADIUS, local; local, RADIUS).

For example, when the Netlogin MAC authentication database order is local, RADIUS, if the authentication of a MAC client fails through a local database, RADIUS is used for authentication. If RADIUS also fails authentication, the client is moved to the authentication failure VLAN.

Authentication service is considered to be unavailable for RADIUS in the following cases:

- RADIUS server is not running.
- RADIUS server is not configured on the switch.
- RADIUS server is configured but not enabled on the switch.



Note

If web is enabled on a port where Dot1x or MAC is also enabled, the authentication failure/service-unavailable VLAN configuration is not applicable to those clients where Dot1x or MAC clients that fail authentication or where authentication service is not available.

Example

The following example adds the service-unavailable VLAN "v1" on tagged ports 1 and 2:

```
# configure netlogin authentication service-unavailable add vlan v1 ports 1,2 tagged
```

History

This command was first available in ExtremeXOS 12.1.

The ability to configure multiple service-unavailable VLANs was added in ExtremeXOS 30.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure netlogin banner

```
configure netlogin banner banner
```

Description

Configures the network login page banner.

Syntax Description

<i>banner</i>	Specifies the HTML code for the banner.
---------------	---

Default

The default banner is the Extreme Networks logo.

Usage Guidelines

The banner is a quoted, HTML string, that will be displayed on the network login page. The string is limited to 1024 characters.

This command applies only to the web-based authentication mode of network login.

Example

The following command configures the network login page banner:

```
configure netlogin banner "<html><head>Please Login</head></html>"
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure netlogin base-url

```
configure netlogin base-url url
```

Description

Configures the base URL for network login.

Syntax Description

<i>url</i>	Specifies the base URL for network login. Note: The netlogin base-url is restricted to 79 characters.
------------	---

Default

The base URL default value is “network-access.com.”

Usage Guidelines

When you login using a web browser, you are redirected to the specified base URL, which is the DNS name for the switch.

You must configure a DNS name of the type “www.xx...xx.xxx” or “xx...xx.xxx”.

This command applies only to the web-based authentication mode of network login.

Example

The following command configures the network login base URL as access.net:

```
configure netlogin base-url access.net
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure netlogin delete mac-list

```
configure netlogin delete mac-list [mac {mask}] | default
```

Description

Deletes an entry from the MAC address list for MAC-based network login.

Syntax Description

<i>mac</i>	Specifies the MAC address to delete.
<i>mask</i>	Specifies the number of bits to use for the mask.
default	Specifies the default entry.

Default

N/A.

Usage Guidelines

Use this command to delete an entry from the MAC address list used for MAC-based network login.

Use this command to remove the default MAC-list configuration after running **enable netlogin mac**.

Example

The following command deletes the MAC address 10:20:30:40:50:60 from the list:

```
configure netlogin delete mac-list 10:20:30:40:50:60
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure netlogin delete proxy-port

```
configure netlogin delete proxy-port tcp_port
```

Description

Configure the ports that are to be hijacked and redirected for HTTP or HTTPS traffic.

Syntax Description

<i>tcp_port</i>	Specifies the port to be hijacked.
-----------------	------------------------------------

Default

N/A.

Usage Guidelines

This command allows you to unconfigure the ports that will be hijacked and redirected for HTTP or HTTPS traffic.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure netlogin dot1x eapol-transmit-version

```
configure netlogin dot1x eapol-transmit-version eapol-version
```

Description

Configures the default EAPOL version sent in transmitted packets for network login.

Syntax Description

<i>eapol-version</i>	Specifies the EAPOL version. Choices are "v1" or "v2".
----------------------	--

Default

The default is "v1".

Usage Guidelines

Although the ExtremeXOS software supports EAPOL version 2, some clients do not yet accept the version 2 EAPOL packets. The packet format for the two versions is the same.

Example

The following command changes the EAPOL version to 2:

```
configure netlogin dot1x eapol-transmit-version v2
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure netlogin dot1x guest-vlan

```
configure netlogin dot1x guest-vlan vlan_name {ports port_list}
```

Description

Configures a guest VLAN for 802.1X authentication network login.

Syntax Description

<i>vlan_name</i>	Specifies the name of the guest VLAN.
<i>port_list</i>	Specifies one or more ports or slots and ports. If the ports keyword is not used, the command applies to all ports.

Default

N/A.

Usage Guidelines

This command configures the guest VLAN for 802.1X on the current virtual router (VR).



Note

Beginning with ExtremeXOS 11.6, you can configure guest VLANs on a per port basis, which allows you to configure more than one guest VLAN per VR. In ExtremeXOS 11.5 and earlier, you can only configure guest VLANs on a per VLAN basis, which allows you to configure only one guest VLAN per VR.

If you do not specify any ports, the guest VLAN is configured for all ports.

Each port can have a different guest VLAN.

A guest VLAN provides limited or restricted network access if a supplicant connected to a port does not respond to the 802.1X authentication requests from the switch. A port always moves untagged into the guest VLAN.

Keep in mind the following when configuring guest VLANs:

- You must create a VLAN and configure it as a guest VLAN before enabling the guest VLAN feature.
- Configure guest VLANs only on network login ports with 802.1X enabled.
- Movement to guest VLANs is not supported on network login ports with MAC-based or web-based authentication.
- 802.1X must be the only authentication method enabled on the port for movement to guest VLAN.

- No supplicant on the port has 802.1X capability.
- You configure only one guest VLAN per virtual router interface.

**Note**

The supplicant does not move to a guest VLAN if it fails authentication after an 802.1X exchange; the supplicant moves to the guest VLAN only if it does not respond to an 802.1X authentication request.

Modifying the Supplicant Timer

By default, the switch attempts to authenticate the supplicant every 30 seconds for a maximum of three tries. If the supplicant does not respond to the authentication requests, the client moves to the guest VLAN. The number of authentication attempts is not a user-configured parameter.

To modify the supplicant response timer, use the following command and specify the `supp-resp-timeout` parameter:

```
configure netlogin dot1x timers [{server-timeout server_timeout}  
{quiet-period quiet_period} {reauth-period reauth_period} {reauth-  
max max_num_reauths}] {supp-resp-timeout supp_resp_timeout}]
```

If a supplicant on a port in the guest VLAN becomes 802.1X-capable, the switch starts processing the 802.1X responses from the supplicant. If the supplicant is successfully authenticated, the port moves from the guest VLAN to the destination VLAN specified by the RADIUS server.

Enabling Guest VLANs

To enable the guest VLAN, use the following command:

```
enable netlogin dot1x guest-vlan ports [all |ports]
```

Example

The following command creates a guest VLAN for 802.1X named guest for all ports:

```
configure netlogin dot1x guest-vlan guest
```

The following command creates a guest VLAN named guest for ports 2 and 3:

```
configure netlogin dot1x guest-vlan guest ports 2,3
```

History

This command was first available in ExtremeXOS 11.2.

The ports option was added in ExtremeXOS 11.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure netlogin dot1x tag-eapol

```
configure netlogin dot1x tag-eapol [on | off]
```

Description

Configures receiving tagged EAPOL packets on dot1x-enabled ports.

Syntax Description

on	Turns EAPOL-tagged frames feature on.
off	Turns EAPOL-tagged frames feature off . Default is off.

Default

Default is off.

Usage Guidelines

When this feature is on and switch receives tagged EAPOL packet on dot1x-enabled ports, tagged EAPOL response is sent out on those ports. On untagged ports, the EAPOL frames are sent untagged. When this feature is off, switch sends untagged EAPOL packets on all the tagged/untagged ports. This command allows you to authenticate dot1x users on tagged and untagged ports.

This command is applicable only for policy-enabled mode.

Example

The following example enables the switch to send tagged EAPOL packets:

```
configure netlogin dot1x tag-eapol on
```

History

This command was first available in ExtremeXOS 22.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure netlogin dot1x timers

```
configure netlogin dot1x timers [{server-timeout server_timeout}  
  {quiet-period quiet_period} {reauth-period reauth_period {reauth-max  
  max_num_reauths}} {reauthentication [on | off] {supp-resp-timeout  
  supp_resp_timeout}}]
```


Description

Configures the 802.1X timers for network login.

Syntax Description

server-timeout	Specifies the timeout period for a response from the <i>RADIUS</i> server. The range is 1 to 120 seconds.
quiet-period	Specifies the time for which the switch will not attempt to communicate with the supplicant after authentication has failed. The range is 0 to 65535 seconds.
reauth-period	Specifies time after which the switch will attempt to re-authenticate an authenticated supplicant. The range is 0 to 86,400 seconds.
reauth-max	Specifies the maximum reauthentication counter value. The range is 1 to 10.
supp-resp-timeout	Specifies the time for which the switch will wait for a response from the supplicant. The range is 1 to 120 seconds.
reauthentication	Enables or disables dot1x reauthentication
on	Enables reauthentication.
off	Disables reauthentication.

Default

The defaults are as follows:

- server-timeout—30 seconds.
- quiet-period—60 seconds.
- reauth-period—3600 seconds.
- reauth-max—3.
- supp-resp-timeout—30 seconds.

Usage Guidelines

To disable re-authentication, specify 0 for the reauth-period parameter. (If reauth-period is set to 0, reauth-max value doesn't apply.)

If you attempt to configure a timer value that is out of range (not supported), the switch displays an error message. The following is a list of sample error messages:

- server-timeout—ERROR: RADIUS server response timeout out of range (1..120 sec)
- quiet-period—%% Invalid number detected at '^' marker. %% Input number must be in the range [0, 65535].
- reauth-period—%% Invalid input detected at '^' marker. %% Input number must be in the range [0, 86400].
- reauth-max—ERROR: Re-authentication counter value out of range (1..10)

- `supp-resp-timeout—ERROR`: Input number must be in the range [1, 10].
- greater than RADIUS timeout—Dot1x server timeout should be configured with a value greater than the RADIUS server timeout.

To display the 802.1X timer settings, use the `show netlogin` command with and without the **dot1x** option.

If reauthentication is enabled by this command, the session-timeout value sent from RADIUS has priority. If no value is sent from RADIUS, then the locally configured `reauth_period` defines the reauthentication period.

If the locally configured value is "0" with reauthentication off, and if any session timeout value sent from RADIUS is ignored, the locally configured "0" takes precedence.

Example

The following command changes the 802.1X server-timeout to 10 seconds:

```
configure netlogin dot1x timers server-timeout 10
```

History

This command was first available in ExtremeXOS 11.1.

The reauth-max keyword was added in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure netlogin dynamic-vlan

```
configure netlogin dynamic-vlan [disable | enable]
```

Description

Configures the switch to automatically and dynamically create a VLAN after receiving authentication requests from one or more supplicants (clients).

Syntax Description

disable	Specifies that the switch does not automatically create dynamic VLANs. This is the default behavior.
enable	Specifies that the switch automatically create dynamic VLANs.

Default

The default is disabled.

Usage Guidelines

Use this command to configure the switch to dynamically create a VLAN. If configured for dynamic VLAN creation, the switch automatically creates a supplicant VLAN that contains both the supplicant's physical port and one or more uplink ports.

A dynamically created VLAN is only a Layer 2 bridging mechanism; this VLAN does not work with routing protocols to forward traffic. After the switch unauthenticates all of the supplicants from the dynamically created VLAN, the switch deletes that VLAN.



Note

Dynamically created VLANs do not support the session refresh feature of web-based network login because dynamically created VLANs do not have an IP address. Also, dynamic VLANs are not supported on ports when *STP* and network login are both configured on the ports.

By dynamically creating and deleting VLANs, you minimize the number of active VLANs configured on your edge switches. In addition, the *RADIUS* server forwards VSA information to dynamically create the VLAN thereby simplifying switch management. A key difference between dynamically created VLANs and other VLANs is that the switch does not save dynamically created VLANs. Even if you use the save command, the switch does not save a dynamically created VLAN.

Supported Vendor Specific Attributes

To prevent conflicts with existing VLANs on the switch, the RADIUS server uses Vendor Specific Attributes (VSAs) to forward VLAN information, including VLAN ID, to the switch. The following list specifies the supported VSAs for configuring dynamic network login VLANs:

- Extreme: Netlogin-VLAN-ID (VSA 209).
- IETF: Tunnel-Private-Group-ID (VSA 81).
- Extreme: Netlogin-Extended-VLAN (VSA 211).



Note

If the ASCII string only contains numbers, it is interpreted as the VLAN ID. Dynamic VLANs only support numerical VLAN IDs; VLAN names are not supported.

The switch automatically generates the VLAN name in the following format: SYS_NLD_ *TAG* where *TAG* specifies the VLAN ID. For example, a dynamic network login VLAN with an ID of 10 has the name SYS_NLD_0010.

Specifying the Uplink Ports

To specify one or more ports as tagged uplink ports that are added to the dynamically created VLAN, use the following command: `configure netlogin dynamic-vlan uplink-ports`

The uplink ports send traffic to and from the supplicants from the core of the network.

By default the setting is none. For more information about this command, see the usage guidelines for `configure netlogin dynamic-vlan uplink-ports`.

Viewing Status Information

To display summary information about all of the VLANs on the switch, including any dynamic VLANs currently operating on the switch, use the following command: `show vlan`

If the switch dynamically creates a VLAN, the VLAN name begins with SYS_NLD_ and the output contains a d flag for the dynamically created VLAN.

To display the status of dynamic VLAN configuration on the switch, use the following command: `show netlogin`

The switch displays the current state of dynamic VLAN creation (enabled or disabled) and the uplink port(s) associated with the dynamic VLAN.

Example

The following example automatically adds ports 1:1-1:2 to the dynamically created VLAN as uplink ports:

```
configure netlogin dynamic-vlan uplink-ports 1:1-1:2
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure netlogin dynamic-vlan uplink-ports

```
configure netlogin dynamic-vlan uplink-ports [port_list | none]
```

Description

Specifies which port(s) are added as tagged, uplink ports to the dynamically created VLANs for network login.

Syntax Description

<i>port_list</i>	Specifies one or more ports to add to the dynamically created VLAN for network login.
none	Specifies that no ports are added. This is the default setting.

Default

The default setting is none.

Usage Guidelines

Use this command to specify which port(s) are used as uplink ports and added to the dynamically created VLAN for network login. The uplink ports send traffic to and from the supplicants from the core of the network.

Uplink ports should not be configured for network login (network login is disabled on uplink ports). If you specify an uplink port with network login enabled, the configuration fails and the switch displays an error message similar to the following:

```
ERROR: The following ports have NetLogin enabled: 1, 2
```

If this occurs, select a port with network login disabled.

Enabling Dynamic Network Login VLANs

To configure the switch to dynamically create a VLAN upon receiving an authentication response, use the following command:

```
configure netlogin dynamic-vlan [disable | enable]
```

By default, the setting is disabled. For more detailed information about this command, see the usage guidelines [configure netlogin dynamic-vlan uplink-ports](#).

Viewing Status Information

To display summary information about all of the VLANs on the switch, including any dynamic VLANs currently operating on the switch, use the following command:

```
show vlan
```

If the switch dynamically creates a VLAN, the VLAN name begins with SYS_NLD_ and the output contains a d flag for the dynamically created VLAN.

To display the status of dynamic VLAN configuration on the switch, use the following command:

```
show netlogin
```

The switch displays the current state of dynamic VLAN creation (enabled or disabled) and the uplink port(s) associated with the dynamic VLAN.

Example

The following command configures the switch to add ports 1:1-1:2 to the dynamically created network login VLAN:

```
configure netlogin dynamic-vlan uplink-ports 1:1-1:2
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure netlogin idle-timeout

```
configure netlogin idle-timeout {convergence-endpoint | dot1x | mac |
web-based} timeout
```

Description

This command clears multiple authentication properties for one or more ports.

Syntax Description

dot1x	IEEE 802.1X Port-based network access control.
mac	MAC authentication.
web-based	Web-based authentication.
convergence-endpoint	Convergence-endpoint authentication.
<i>timeout</i>	Number of seconds before idle timeout (range 0-172800).

Default

Timeout = 300 seconds.

Usage Guidelines

This command appears in `show configuration {module-name} {detail}` for "policy" rather than "netlogin."

History

This command was first available in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure netlogin local-user security-profile

```
configure netlogin local-user user-name security-profile
security_profile
```

Description

Changes a previously associated security profile.

Syntax Description

<i>user-name</i>	Specifies the name of an existing local network login account.
<i>security_profile</i>	Specifies a security profile string during account creation.

Default

N/A.

Usage Guidelines

Use this command to change any previously associated security profiles on the switch.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure netlogin local-user

```
configure netlogin local-user user-name {vlan-vsa [[{tagged | untagged}  
[vlan_name | vlan_tag]] | none]}
```

Description

Configures an existing local network login account.

Syntax Description

<i>user-name</i>	Specifies the name of an existing local network login account.
tagged	Specifies that the client be added as tagged.
untagged	Specifies that the client be added as untagged.
<i>vlan_name</i>	Specifies the name of the destination <u>VLAN</u> .
<i>vlan_tag</i>	Specifies the VLAN ID, tag, of the destination VLAN.
none	Specifies that the VSA 211 wildcard (*) is applied, only if you do not specify tagged or untagged.

Default

N/A.

Usage Guidelines

Use this command to modify the attributes of an existing local network login account. You can update the following attributes associated with a local network login account:

- Password of the local network login account.
- Destination VLAN attributes including: adding clients tagged or untagged, the name of the VLAN, and the VLAN ID.



Note

Passwords are case-sensitive and must have a minimum of 1 character and a maximum of 32 characters.

You must create a local network login account before using this command. To create a local network login user name and password, use the following command:

```
create netlogin local-user user-name {encrypted} {password} {vlan-  
vsa [[{tagged | untagged} [vlan_name] | vlan_tag]]} {security-  
profilesecurity_profile}
```

If the switch displays a message similar to the following:

```
* Switch # configure netlogin local-user purplenet  
^  
%% Invalid input detected at '^' marker.
```

You might be attempting to modify a local network login account that is not present on the switch, or you might have incorrectly entered the account name. To confirm the names of the local network login accounts on your switch, use the following command:

```
show netlogin local-users
```

Additional Requirements

This command applies only to the web-based and MAC-based modes of network login. 802.1X network login does not support local database authentication.

You must have administrator privileges to use this command. If you do not have administrator privileges, the switch displays a message similar to the following:

```
This user does not have permissions for this command.
```

Passwords are case-sensitive. Passwords must have a minimum of 0 characters and a maximum of 32 characters. If you attempt to create a password with more than 32 characters, the switch displays the following message after you re-enter the password:

```
Password cannot exceed 32 characters
```

Example

This section contains the following examples:

- Updating the password.
- Modifying destination VLAN attributes.

Updating the Password

The following command updates the password of an existing local network login account:

```
configure netlogin local-user megtest
```

After you enter the local network login user name, press [Enter]. The switch prompts you to enter a password; however, the switch does not display the password. At the prompt enter the new password:

```
password:
```

After you enter the new password, press [Enter]. The switch then prompts you to re-enter the password:

```
Reenter password:
```

Updating VLAN Attributes

You can add a destination VLAN, change the destination VLAN, or remove the destination from an existing local network login account. This example changes the destination VLAN for the specified local network login account:

```
configure netlogin local-user megtest vlan-vsa green
```

History

This command was first available in ExtremeXOS 11.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure netlogin mac timers reauth-period

```
configure netlogin mac ports [port_list | all] timers [{reauth-period [reauth_period]} {reauthentication [on|off]} {delay [delay_period]}]
```

Description

Configures the reauthentication period for network login MAC-based authentication.

Syntax Description

reauth_period	Specifies time after which the switch will attempt to re-authenticate an authenticated supplicant. The range is 0, 30 to 86,400 seconds.
reauthentication	Configure mac reauthentication.
on	MAC reauthentication is enabled.
off	MAC reauthentication is disabled.

delay	Configure MAC authentication delay period.
<i>delay_period</i>	MAC authentication delay period. 0-120 seconds range.

Default

The default is 0 (disabled).

Usage Guidelines

This command allows you to configure the reauth-period for network login MAC-based authentication. The session-timeout configuration on the *RADIUS* server overrides the reauth-period if it has been configured.

In MAC mode, if reauthentication is turned off, globally and per-port, using this command, a session timeout sent by RADIUS takes precedence and local timers are ignored.

Example

The following command configures a MAC authentication delay period of 100 seconds on port 39:

```
configure netlogin mac ports 39 timers delay 100
```

History

This command was first available in ExtremeXOS 12.1.

The **delay** keyword and variable were added in ExtremeXOS 21.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure netlogin mac username case

```
configure netlogin mac username case (lower | upper)
```

Description

Sets option to send the Network Login (NetLogin) MAC Authentication MAC address in either uppercase or lowercase for user name or password.

Syntax Description

netlogin	Configures NetLogin specific settings.
mac	Configures NetLogin settings specific to MAC.
username	Configures MAC user name credential attributes.
case	Configures MAC user name case.

lower	Use lowercase (for example: aa:bb:cc:dd:ee:ff).
upper	Use uppercase (for example: AA:BB:CC:DD:EE:FF). Default.

Default

By default, the uppercase is used.

Usage Guidelines

When the user name case is configured as lowercase, if the client with MAC address aa:bb:cc:dd:ee:ff sends a frame, Netlogin MAC sends “aabbccddeeff” (default “None” delimiter) as username and default password for authentication.

Example

The following example sets the NetLogin MAC to be sent in lowercase:

```
# configure netlogin mac username case lower
```

History

This command was first available in ExtremeXOS 30.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure netlogin mac username format

```
configure netlogin mac username format [hyphenated | colon-separated | none]
```

Description

Configures the NetLogin MAC username format used when sending out for authentication to a RADIUS server.

Syntax Description

mac	Configure Network Login settings specific to MAC.
username	Configure username credential attributes.
format	Configure username format.
hyphenated	Hyphen separator (XX-XX-XX-XX-XX-XX).
colon-separated	Colon separator (XX:XX:XX:XX:XX:XX).
none	No separator (XXXXXXXXXXXX) (This is the default).

Default

No separator is the default.

Example

The following example sets the MAC username format with colon separator:

```
configure netlogin mac username format colon-separated
```

History

This command was first available in ExtremeXOS 12.1.

The **colon-separated** option was added in ExtremeXOS 22.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure netlogin move-fail-action

```
configure netlogin move-fail-action [authenticate | deny]
```

Description

Configures the action network login takes if a VLAN move fails. This can occur if two clients attempt to move to an untagged VLAN on the same port.

Syntax Description

authenticate	Specifies that the client is authenticated.
deny	Specifies that the client is not authenticated. This is the default setting.

Default

The default setting is deny.

Usage Guidelines

Use this command to specify how network login behaves if a VLAN move fails. Network login can either authenticate the client on the current VLAN or deny the client.

The following describes the parameters of this command if two clients want to move to a different untagged VLAN on the same port:

- **authenticate**—Network login authenticates the first client that requests a move and moves that client to the requested VLAN. Network login authenticates the second client but does not move that client to the requested VLAN. The second client moves to the first client’s authenticated VLAN.

- deny—Network login authenticates the first client that requests a move and moves that client. Network login does not authenticate the second client.

To view the current move-fail-action setting on the switch, use the `show netlogin` command.

Example

The following command configures network login to authenticate the client on the current VLAN:

```
configure netlogin move-fail-action authenticate
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure netlogin port allow egress-traffic

```
configure netlogin ports [port_list | all] allow egress-traffic [none |  
unicast | broadcast | all_cast]
```

Description

Configures the egress traffic in an unauthenticated state.

Syntax Description

all	Specifies all network login ports.
<i>port_list</i>	Specifies one or more network login ports.
none	Specifies that no traffic is sent out if if no authenticated clients exist on the <u>VLAN</u> .
unicast	Specifies that the unicast flooding traffic for the VLANs on the network login enabled port be sent.
broadcast	Specifies that the broadcast traffic for the VLANs on the network login enabled port be sent.
all_cast	Specifies that the broadcast and unicast flooding traffic for the VLANs on the network login enabled port be sent.

Default

The default is none.

Usage Guidelines

This command allows you to configure the egress traffic in an unauthenticated state on a per-port basis.

Enabling ONEPolicy removes the action of this command. This command is supported only in non-policy mode

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure netlogin ports

```
configure netlogin ports [all | port_list] [allowed-users allowed_users
| authentication mode [optional | required] | trap [all-traps | no-
traps | [{success} {failed} {terminated} {max-reached}]]]
```

Description

Use this command to set the [NetLogin](#) trap setting for ports.

Syntax Description

all	Configure all ports in the system.
<i>port_list</i>	List of ports to configure.
allowed-users	Number of users allowed per port. Only applicable if the ONEPolicy feature is enabled.
<i>allowed_users</i>	Number of users allowed per port.
authentication mode	Port authentication mode. Only applicable if the ONEPolicy feature is enabled.
optional	Authentication optional. Only applicable if ONEPolicy is enabled.
required	Authentication required. Only applicable if ONEPolicy is enabled.
all-traps	Enable sending all trap types. Only applicable if the ONEPolicy feature is enabled.
no-traps	Disable sending all trap types. Only applicable if the ONEPolicy feature is enabled.
success	Enable sending success trap.
fails	Enable sending failed trap.
terminated	Enable sending terminated trap.
max-reached	Enable sending max number users reached trap. This is applicable in ONEPolicy mode only.

Default

By default, all traps are sent in both ONEPolicy mode and non-ONEPolicy mode.

Usage Guidelines

The following command options are only applicable if ONEPolicy is enabled. They have no effect without ONEPolicy being enabled:

- **authentication mode** [**optional** | **required**]
- **allowed-users** *allowed_users*
- **all-traps** | **no-traps** | [{**success**} {**failed**} {**terminated**} {**max-reached**}]

This command appears in `show configuration {module-name} {detail}` for "policy" and "netlogin."

The **no-traps** configuration is retained after save and reboot.

Trap configurations after applying **no-traps** are appended until **no-traps** is configured again (for example: **no-traps** configuration followed by **success**, and then **terminated** traps, sends success and terminated traps:

```
# configure netlogin ports 1 trap no-traps
# show configuration "policy"
  **no traps commands appear due to no-traps being configured
# configure netlogin ports 1 trap success
# show configuration "policy"
# Module policy configuration.
# configure netlogin ports 1 trap success
  **success traps command appears
# configure netlogin ports 1 trap terminated
# show configuration "policy"
# Module policy configuration.
# configure netlogin ports 1 trap success
# configure netlogin ports 1 trap terminated
  **success and terminated traps commands appear
```

Example

This example shows how to enable all NetLogin port trap setting:

```
configure netlogin trap port 1:1 all
```

History

This command was first available in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure netlogin ports mode

```
configure netlogin ports [all | port_list] mode [mac-based-vlans | port-based-vlans]
```

Description

Configures the network login port's mode of operation.

Syntax Description

all	Specifies all netlogin ports.
<i>port_list</i>	Specifies one or more network login ports.
mac-based-vlans	Allows more than one untagged <u>VLAN</u> .
port-based-vlans	Allows only one untagged VLAN. This is the default behavior.

Default

The default setting is **port-based-vlans**.

Usage Guidelines

Use this command to configure network login MAC-based VLANs on a network login port.

If you modify the mode of operation to mac-based-vlans and later disable all network login protocols on that port, the mode of operation automatically returns to port-based-vlans.

When you change the network login port's mode of operation, the switch deletes all currently known supplicants from the port and restores all VLANs associated with that port to their original state. In addition, by selecting mac-based-vlans, you are unable to manually add or delete untagged VLANs from this port. Network login now controls these VLANs.

With network login MAC-based operation, every authenticated client has an additional FDB flag that indicates a translation MAC address. If the supplicant's requested VLAN does not exist on the port, the switch adds the requested VLAN.

Configuration of **port-based-vlans** is lost if ONEPolicy is enabled.

Important Rules and Restrictions

This section summarizes the rules and restrictions for configuring network login MAC-based VLANs:

- If you attempt to configure the port's mode of operation before enabling network login, the switch displays an error message similar to the following:
ERROR: The following ports do not have NetLogin enabled; 1

To enable network login on the switch, use the following command to enable network login and to specify an authentication method (for example, 802.1X—identified as dot1.x in the CLI):


```
enable netlogin dot1x
```

To enable network login on the ports, use the following command to enable network login and to specify an authentication method (for example, 802.1X—identified as dot1x in the CLI):

```
enable netlogin ports 1:1 dot1x
```

- On ExtremeXOS versions prior to 12.0 on switches other than the ExtremeSwitching series switches, 10 Gigabit Ethernet ports such as those on the uplink ports on the switches do not support network login MAC-based VLANs.

If you attempt to configure network login MAC-based VLANs on 10 Gigabit Ethernet ports, the switch displays an error message similar to the following:

```
ERROR: The following ports do not support the MAC-Based VLAN mode; 1, 2, 10
```

- You can have a maximum of 1,024 MAC addresses per ExtremeSwitching switch.

Displaying FDB Information

To view network login-related FDB entries, use the following command:

```
show fdb netlogin [all | mac-based-vlans]
```

The following is sample output from the `show fdb netlogin mac-based-vlans` command:

Mac	Vlan	Age	Use	Flags	Port List
00:04:96:10:51:80	VLONE(0021)	0086	0000	n m	v 1:11
00:04:96:10:51:81	VL TWO(0051)	0100	0000	n m	v 1:11
00:04:96:10:51:91	VL TWO(0051)	0100	0000	n m	v 1:11

Flags : d - Dynamic, s - Static, p - Permanent, n - NetLogin, m - MAC, i - IP,
 x - IPX, l - lockdown MAC, M - Mirror, B - Egress Blackhole,
 b - Ingress Blackhole, v - NetLogin MAC-Based VLAN.

The flags associated with network login include:

- v—Indicates the FDB entry was added because the port is part of a MAC-based virtual port/VLAN combination.
- n—Indicates the FDB entry was added by network login.

Displaying Port and VLAN Information

To view information about the VLANs that are temporarily added in MAC-based mode for network login, use the following command:

```
show ports port_list information detail
```

The following is sample output from this command:

```
Port: 1
Virtual-router: VR-Default
Type: UTP
Random Early drop: Disabled
Admin state: Enabled with auto-speed sensing auto-duplex
Link State: Active, 100Mbps, full-duplex
Link Counter: Up 1 time(s)
VLAN cfg:
```

```
Name: Default, Internal Tag = 1(MAC-Based), MAC-limit = No-limit
...<truncated output>
Egress 802.1p Replacement:      Disabled
NetLogin:                      Enabled
NetLogin authentication mode:   Mac based
NetLogin port mode:            MAC based VLANs
Smart redundancy:              Enabled
Software redundant port:       Disabled
auto-polarity:                 Enabled
```

The added output displays information about the mode of operation for the network login port.

- VLAN cfg—The term MAC-based appears next to the tag number.
- *NetLogin* port mode—This output was added to display the port mode of operation. Mac based appears as the network login port mode of operation.

To view information about the ports that are temporarily added in MAC-based mode for network login, due to discovered MAC addresses, use the following command:

```
show vlan detail
```

The following is sample output from this command:

```
VLAN Interface with name Default created by user
Tagging:      802.1Q Tag 1
Priority:      802.1P Priority 0
Virtual router: VR-Default
STPD:         s0(Disabled,Auto-bind)
Protocol:     Match all unfiltered protocols
Loopback:     Disable
NetLogin:     Disabled
Rate Shape:   Disabled
QosProfile:   None configured
Ports: 26.    (Number of active ports=2)
Untag:  *1um, *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
Flags: (*) Active, (!) Disabled, (g) Load Sharing port
(b) Port blocked on the vlan, (a) Authenticated NetLogin Port
(u) Unauthenticated NetLogin port, (m) Mac-Based port
```

The flags associated with network login include:

- a—Indicates an authenticated network login port.
- u—Indicates an unauthenticated network login port.
- m—Indicates that the network login port operates in MAC-based mode.

Example

The following command configures the network login ports mode of operation:

```
configure netlogin ports 1:1-1:10 mode mac-based-vlans
```

History

This command was first available in ExtremeXOS 11.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure netlogin ports no-restart

```
configure netlogin ports [all | port_list] no-restart
```

Description

Disables the network login port restart feature.

Syntax Description

all	Specifies all network login ports.
<i>port_list</i>	Specifies one or more network login ports.

Default

The default setting is no-restart; the network login port restart feature is disabled.

Usage Guidelines

Use this command to disable the network login port restart feature on a network login port.

Configure network login port restart on ports with directly attached supplicants. If you use a hub to connect multiple supplicants, only the last unauthenticated supplicant causes the port to restart.

Enabling ONEPolicy removes the action of this command. This command is supported only in non-policy mode

Displaying the Port Restart Configuration

To display the network login settings on the port, including the configuration for port restart, use the following command:

```
show netlogin port port_list
```

Output from this command includes the enable/disable state for network login port restart.

Example

The following command disables network login port restart on port 1:1:

```
configure netlogin ports 1:1 no-restart
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure netlogin ports restart

```
configure netlogin ports [all | port_list] restart
```

Description

Enables the network login port restart feature.

Syntax Description

all	Specifies all network login ports.
<i>port_list</i>	Specifies one or more network login ports.

Default

The default setting is no-restart; the network login port restart feature is disabled.

Usage Guidelines

Use this command to enable the network login port restart feature on a network login port. This allows network login to restart specific network login-enabled ports when the last authenticated supplicant releases, regardless of the configured protocols on the port.

Configure network login port restart on ports with directly attached supplicants. If you use a hub to connect multiple supplicants, only the last unauthenticated supplicant causes the port to restart.

Enabling ONEPolicy removes the action of this command. This command is supported only in non-policy mode

Displaying the Port Restart Configuration

To display the network login settings on the port, including the configuration for port restart, use the following command:

```
show netlogin port port_list
```

Output from this command includes the enable/disable state for network login port restart.

Example

The following command enables network login port restart on port 1:1:

```
configure netlogin ports 1:1 restart
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure netlogin redirect-page

```
configure netlogin redirect-page url
```

Description

Configures the redirect URL for Network Login.

Syntax Description

<i>url</i>	Specifies the redirect URL for Network Login.
------------	---

Default

The redirect URL default value is "http://www.extremenetworks.com"; the default port value is 80.

Usage Guidelines

In ISP mode, you can configure network login to be redirected to a base page after successful login using this command. If a *RADIUS* server is used for authentication, then base page redirection configured on the RADIUS server takes priority over this configuration.

You must configure a complete URL starting with http:// or https://

You can also configure a specific port location at a specific target URL location. For example, you can configure a target port 8080 at extremenetworks.com with the following command:

```
configure netlogin redirect-page "www.extremenetworks.com:8080"
```

This command applies only to the web-based authentication mode of Network Login.

Example

The following command configures the redirect URL as http://www.extremenetworks.com/support:

```
configure netlogin redirect-page http://www.extremenetworks.com/support
```

History

This command was first available in ExtremeXOS 11.1.

Support for HTTPS was introduced in ExtremeXOS 11.2.

Target port support was introduced in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure netlogin session-refresh

```
configure netlogin session-refresh {refresh_seconds}
```

Description

Configures network login session refresh.

Syntax Description

<i>refresh_seconds</i>	Specifies the session refresh time for network login in seconds.
------------------------	--

Default

Enabled, with a value of 180 seconds for session refresh.

Usage Guidelines

Network login sessions can refresh themselves after a configured timeout. After the user has been logged in successfully, a logout window opens which can be used to close the connection by clicking on the Logout link. Any abnormal closing of this window is detected on the switch and the user is logged out after a time interval as configured for session refresh. The session refresh is enabled and set to 360 seconds by default. The value can range from 1 to 3600 seconds. When you configure the network login session refresh for the logout window, ensure that the *FDB* aging timer is greater than the network login session refresh timer.

This command applies only to the web-based authentication mode of network login.

Example

The following command enables network login session refresh and sets the refresh time to 100 seconds:

```
configure netlogin session-refresh 100
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure netlogin session-timeout

```
configure netlogin session-timeout {dot1x | mac | web-based |  
convergence-endpoint} timeout
```

Description

Use this command to set the maximum number of seconds an authenticated session may last before termination of the session.

Syntax Description

dot1x	IEEE 802.1X Port-based network access control.
mac	MAC authentication.
web-based	Web-based authentication.
convergence-endpoint	Convergence-endpoint authentication.
<i>timeout</i>	Number of seconds before session timeout (range 0-172800).

Default

0 seconds.

Usage Guidelines

A value of zero may be superseded by a session timeout value provided by the authenticating server. For example, if a session is authenticated by a *RADIUS* server, that server may encode a session-timeout attribute in its authentication response.

The specifications from this command appear in `show configuration {module-name} {detail}` for "policy" and "netlogin."

If you want to scale to 65,000 authenticated users, use a session timeout value of at least 300 minutes.

Example

The following example shows how to set the session-timeout value for an active session, for mac authentication to 500 seconds:

```
configure netlogin session-timeout mac 500
```

History

This command was first available in ExtremeXOS 16.1.

The **convergence-endpoint** option was added in ExtremeXOS 22.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure netlogin trap

```
configure netlogin trap max-users [enable | disable]
```

Description

Use this command to set the *NetLogin* system traps.

Syntax Description

enable	Enable sending traps when max users reached in system
disable	Disable sending traps when max users reached in system

Default

Disabled.

Usage Guidelines

The specifications from this command appear in `show configuration {module-name} {detail}` for "policy" and "netlogin."

Example

This example shows how to enable the NetLogin maximum users trap setting:

```
configure netlogin trap max-users enabled
```

History

This command was first available in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure netlogin vlan

```
configure netlogin vlan vlan_name
```

Description

Configures the *VLAN* for Network Login.

Syntax Description

vlan	Specifies the VLAN for Network Login.
-------------	---------------------------------------

Default

N/A.

Usage Guidelines

This command will configure the VLAN used for unauthenticated clients. One VLAN needs to be configured per VR. To change the VLAN, network login needs to be disabled. Network login can only be enabled when a VLAN is assigned (and no ports are configured for it).

By default no VLAN is assigned for network login.

Example

The following command configures the VLAN login as the network login VLAN:

```
configure netlogin vlan login
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure network-clock gptp bmca

```
configure network-clock gptp bmca [ on | off ]
```

Description

This command configures the Best Master Clock Algorithm (BMCA) as part of gPTP.

Syntax Description

network-clock	Network Clock
gptp	IEEE 802.1AS Generalized Precision Time Protocol
bmca	Best Master Clock Algorithm
on	Use BMCA to dynamically port roles.
off	Disable BMCA and statically set port roles.

Default

On.

Usage Guidelines

Use this command to configure the BMCA as part of gPTP.

Example

The following example displays output from the show command with BMCA.

```
# show network-clock gptp
gPTP status      : Enabled
BMCA             : [On | Off]
Static slave port : 5 (used when BMCA Off)
--or--
Static slave port : None (used when BMCA Off)
gPTP enabled ports : *1m      *21d  *22d  *47d
Flags:            (*) Active, (!) Administratively disabled,
                  (d) Disabled gPTP port role, (m) Master gPTP port role,
                  (p) Passive gPTP port role,  (s) Slave gPTP port role
```

History

This command was first available in ExtremeXOS 15.7.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure network-clock gptp default-set

```
configure network-clock gptp default-set [{priority1 priority1_value}
                                           {priority2 priority2_value}]
```

Description

This command configures the switch's default-set parameters, specifically its grandmaster clock priority values that are used to elect the grandmaster clock in the network.

Syntax Description

<i>priority1_value</i>	The switch's grandmaster clock priority1 value. This is the most significant parameter used to select the grandmaster clock in the network. Lower values indicate higher priority, and 255 prevents the switch from becoming the grandmaster clock.
<i>priority2_value</i>	The switch's grandmaster clock priority2 value. This is one of the least significant parameters used to select the grandmaster clock in the network. Lower values indicate higher priority.

Default

- Priority1_value = 246 (from 802.1AS 8.6.2.1)
- Priority2_value = 248 (from 802.1AS 8.6.2.5)

Usage Guidelines

Use this command to configure the switch's default-set parameters, specifically its grandmaster clock priority values that are used to elect the grandmaster clock in the network. The Best Master Clock Algorithm uses six parameters from each time-aware system in the network to select the grandmaster clock in the network. Priority1 is the highest precedence value; it allows users to preemptively configure which systems they prefer to be the grandmaster clock. Priority2 is a lower precedence value; it allows users to configure tiebreaker priorities.

The default priority1 values defined by IEEE 802.1AS-2011 clause 8.6.2.1 give preference to network infrastructure systems such as Extreme switches.

Example

```
configure network-clock gtp default-set priority1 248
configure network-clock gtp default-set priority2 100
configure network-clock gtp default-set priority1 248 priority2 100
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure network-clock gtp ports announce

```
configure network-clock gtp ports [port_list {only} | all] announce
[initial-interval log_2_interval | receipt-timeout timeout_count]
```

Description

Configures gTP Announce parameters on the specified ports. Announce messages are used to elect the grandmaster clock and determine the time-synchronous spanning tree.

Syntax Description

<i>port_list</i>	Specifies one or more of the switch's physical ports.
only	Apply change only to specified port, even if port is master of a load sharing group.
all	Specifies all of the switch's physical ports.

<i>log_2_interval</i>	The interval between Announce messages used by the switch on the port when the port is initialized or when the switch receives a message interval request TLV with announceInterval value 126. This value is in log 2 seconds. The valid range of values is -3 ($2^{-3} = 0.125$ seconds) to 17 ($2^{17} = 131072$ seconds).
<i>timeout_count</i>	On a gPTP slave port, the number of announce intervals to wait without receiving an Announce message before assuming the master is no longer sending Announce messages.

Default

- `log_2_interval` = 0 (1 second; 802.1AS-2011 10.6.2.2)
- `timeout_count` = 3 (802.1AS-2011 10.6.3.2)

Usage Guidelines

Use this command to configure gPTP Announce parameters on the specified ports. Announce messages are used to elect the grandmaster clock and determine the time-synchronous spanning tree. Announce selects the grandmaster in the network and establishes the tree from the grandmaster to all other time-aware systems in the network.

initial-interval corresponds to 802.1AS parameter `initialLogAnnounceInterval`.

receipt-timeout corresponds to 802.1AS parameter `announceReceiptTimeout`.

Example

```
# configure network-clock gptp ports 1-2 announce initial-interval 127
# configure network-clock gptp ports all announce receipt-timeout 5
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the AVB feature pack license and have it installed. For complete information about software licensing, including how to obtain and upgrade your license, and which platforms support the AVB feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure network-clock gptp ports peer-delay

```
configure network-clock gptp ports [port_list {only} | all] peer-
delay [{allowed-lost-responses lost_responses_value} {initial-req-
interval log_2_interval} [{asymmetr asymmetry_time [nanoseconds |
microseconds | milliseconds | seconds] | neighbor-thresh [auto |
```

```
neighbor_thresh_time [nanoseconds | microseconds | milliseconds |
seconds]]]{correction-field fractional-ns-only on_off}]
```

Description

Configures gPTP peer delay parameters on the specified ports.

Syntax Description

<i>port_list</i>	Specifies one or more of the switch's physical ports.
only	Apply change only to specified port, even if port is master of a load sharing group.
all	Specifies all of the switch's physical ports.
<i>lost_responses_value</i>	The number of consecutive Peer Delay RequestPdelay_Req messages that the switch must send on a port without receiving a valid response before it considers the port not to be exchanging Peer Delay messages with its neighbor.
<i>log_2_interval</i>	The interval between Peer Delay RequestPdelay_Req messages sent by the switch on the port when the port is initialized or when the switch receives on the port a message interval request TLV with linkDelayInterval value of 126. This value is in log2 seconds. The valid range of values is -3 ($2^{-3} = 0.125$ seconds) to 17 ($2^{17} = 131072$ seconds).
<i>asymmetry_time</i>	The time that the propagation delay from this switch to the neighbor is less than the estimated one-way propagation delay between the switch and its neighbor (which is also the time that the propagation delay from the neighbor to this switch is greater than the estimate). This value is negative if the propagation delay to the neighbor is greater than the estimate. It can be in nanoseconds, microseconds, milliseconds, or seconds. The maximum value is 4,294,967,295 nanoseconds (approximately 4.3 seconds). Let tIR be the propagation delay from this switch (initiator) to the neighbor (responder), tRI be the propagation delay from the neighbor to this switch, and meanPathDelay be the estimated one-way propagation delay. Then: <ul style="list-style-type: none"> • $\text{meanPathDelay} = (\text{tIR} + \text{tRI}) / 2$ • $\text{tIR} = \text{meanPathDelay} - \text{asymmetry_time}$ • $\text{tRI} = \text{meanPathDelay} + \text{asymmetry_time}$
<i>neighbor_thresh_time</i>	The maximum measured mean of the propagation delay between this switch and the neighbor above which the switch considers the port unable to run gPTP. This value can be in nanoseconds, microseconds, milliseconds, or seconds.

auto	Use a media specific default value for the neighbor_thresh_time: <ul style="list-style-type: none"> Copper: 800 nanoseconds. This category includes short range copper cables such as SFP+ Direct Attach and QSRP+ Passive Copper. Multi-mode fiber: 11 microseconds. This category includes the QSFP+ Active Optical cables. 11 microseconds allows 10 microseconds for 100BASE-FX 2 km plus 10% tolerance.) Single-mode fiber: 550 microseconds. This allows 500 microseconds for our “LX100” transceiver plus 10% tolerance. <p>Note: These values may change. A draft of the 802.1AS corrigendum (P802.1AS-Cor-1/D1.1) specifies 800 ns for 100BASE-TX and 1000BASE-T.</p>
correction-field	Specifies configuring the correction field of peer delay messages.
fractional-ns-only	Specifies considering only the fractional nano-second portion for peer delay calculations.
<i>on_off</i>	Consider only fractional nano-second portion, on or off. Default is off.

Default

- Lost_responses_value = 3 (802.1AS 11.5.3)
- Log_2_interval = 0 (1 second; not specified in 802.1AS)
- Asymmetry_time = 0 (802.1AS 10.2.4.8)
- Neighbor_thresh_time = Copper media: 800 nanoseconds, fiber media: 4,294,967,295 nanoseconds
- Considering only the fractional nano-second portion of correction field of peer delay messages if off.

Usage Guidelines

Peer Delay messages determine whether a neighboring system is gPTP capable and measure the propagation delay on the link between the switch and a neighboring gPTP capable system.

- **allowed-lost-responses** corresponds to 802.1AS parameter allowedLostResponses.
- **initial-req-interval** corresponds to 802.1AS parameter initialLogPdelayReqInterval.
- **asymmetry** corresponds to 802.1AS parameter delayAsymmetry.
- **neighbor-thresh** corresponds to 802.1AS parameter neighborPropDelayThresh.

Example

```
configure network-clock gtp ports 1-3 peer-delay allowed-lost-responses 5
configure network-clock gtp ports 1-2 peer-delay initial-log-interval -3
configure network-clock gtp ports 1-2 peer-delay neighbor-thresh 3 nanoseconds
```

History

This command was first available in ExtremeXOS 15.3.

Options to control whether or not you consider only the fractional nano-second portion of correction field of peer delay messages was added in ExtremeXOS 31.1.

Platform Availability

This command is available on platforms that support the AVB feature pack license and have it installed. For complete information about software licensing, including how to obtain and upgrade your license, and which platforms support the AVB feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure network-clock gtp ports sync

```
configure network-clock gtp ports [port_list {only} | all] sync
    [initial-interval log_2_interval receipt-timeout timeout_count]
```

Description

Configures gTP synchronization parameters on the specified ports.

Syntax Description

<i>port_list</i>	Specifies one or more of the switch's physical ports.
only	Apply change only to specified port, even if port is master of a load sharing group.
all	Specifies all of the switch's physical ports.
<i>log_2_interval</i>	The interval between Sync messages used by the switch for the port when the port is initialized or when the switch receives a message interval request TLV with timeSyncInterval value of 126. This value is in log2 seconds. The valid range of values is -3 (2 ⁻³ = 0.125 seconds) to 17 (2 ¹⁷ = 131072 seconds).
<i>timeout_count</i>	On a gTP slave port, the number of sync intervals to wait without receiving a Sync message before assuming the adjacent master port is no longer sending Sync messages.

Default

- *log_2_interval* = -3 (0.125 second; 802.1AS 11.5.2.3)
- *timeout_count* = 3 (802.1AS 10.6.3.1)

Usage Guidelines

Synchronization distributes the time from the grandmaster to all other time-aware systems in the networks.

initial-interval corresponds to 802.1AS parameter `initialLogSyncInterval`.

receipt-timeout corresponds to 802.1AS parameter `syncReceiptTimeout`.

Example

```
configure network-clock gtp ports 1-2 sync initial-interval -1
configure network-clock gtp ports all sync receipt-timeout 5
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the AVB feature pack license and have it installed. For complete information about software licensing, including how to obtain and upgrade your license, and which platforms support the AVB feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure network-clock gtp slave-port

```
configure network-clock gtp slave-port [ port_no | none ]
```

Description

This command allows to you configure the port that will be the slave-port when BMCA is off. All other enabled network gPTP ports will be master ports.

Syntax Description

network-clock	Network Clock
gtp	Variable description, available options, and notes.
slave-port	Configure slave port when Best Master Clock Algorithm is off.
port_no	Port Number of slave port
none	This switch is the Grand Master Clock (GMC).

Default

N/A.

Usage Guidelines

Use this command to you configure the port that will be the slave-port when BMCA is off. All other enabled network gPTP ports will be master ports.

Example

The following example shows the output of the `show network-clock gtp` command with BMCA.

```
gPTP status           : Enabled
```



```

BMCA                : [On | Off]
Static slave port   : 5 (used when BMCA Off)
--or--
Static slave port   : None (used when BMCA Off)
gPTP enabled ports  : *1m      *21d    *22d    *47d
Flags:              (*) Active, (!) Administratively disabled,
                    (d) Disabled gPTP port role, (m) Master gPTP port role,
                    (p) Passive gPTP port role,  (s) Slave gPTP port role

```

History

This command was first available in ExtremeXOS 15.7.

Platform Availability

This command is available on platforms that support the AVB feature pack license and have it installed. For complete information about software licensing, including how to obtain and upgrade your license, and which platforms support the AVB feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure nodealias ports

```
configure nodealias ports [port_list | all] maxentries entries
```

Description

This command modifies the per-port maximum number of alias entries in the Node Alias database. Node Alias discovers information about the end systems on a per-port basis. Information from packets from end systems, such as VLANID, source MAC address, source IP address, protocol, etc. are captured in a database that can be queried.

Syntax Description

nodealias	Node Alias feature that maps source IP address, MAC address, host name, and protocol on a per port basis.
ports	Designates that specified ports should have the specified maximum number of alias entries applied.
<i>port_list</i>	Lists the ports to apply the specified maximum number of alias entries to. Designated as a port list separated by comma (,) or dash (-).
all	Specifies that all ports have the specified maximum number of alias entries.
maxentries	Designates a maximum number of alias entries per port.
<i>entries</i>	The value for the maximum number of aliases entries. The default is 8,192 divided evenly by the number of ports in the switch.

Default

If no value is specified for the maximum number of alias entries, the default is 8,192 divided evenly by the number of ports in the switch.

Usage Guidelines

The per-port limit can be set up to 8,192 for all switch ports. For example, if the switch has 32 ports, you can configure the maximum limit as $32 \times 8,192$. However, the switch can only hold a maximum of 8,192 alias entries per slot.

As a result of snooping one frame, the Node Alias feature may create additional entries to facilitate the searching based on finer details, such as protocol type. For example, when a *BGP* frame is received, two entries are created: one entry with protocol type IP, and another entry with protocol type BGP.

If you change the maximum alias entries to a value that is less than the number entries in the database, the more recent entries are retained.

Example

The following example specifies a maximum of 100 alias entries on all ports:

```
configure nodealias ports all maxentries 100
```

History

This command was first available in ExtremeXOS 22.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ntp key trusted/not-trusted

```
configure ntp key keyid [trusted | not-trusted]
```

Description

Specifies whether an NTP key is trusted or not trusted.

Syntax Description

<i>keyid</i>	Specifies the key ID as a value from 1 to 65534.
trusted	Specifies that the key is in trusted status. To use a specific key for an NTP session, set the key to trusted status.
not-trusted	Specifies that the key is in not trusted status.

Default

An NTP key is not trusted by default.

Usage Guidelines

After an NTP key is created, the generated key is not-trusted by default. To use a specific key for an NTP session, the key must be trusted. The trusted option changes the key to trusted status. The not-trusted option changes the key to untrusted status.

Example

The following command changes NTP key 1 to trusted status:

```
configure ntp key 1 trusted
```

History

This command was first available in ExtremeXOS 12.7.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ntp local-clock none

```
configure ntp local-clock none
```

Description

Removes the internal local clock from the clock source list.

Syntax Description

N/A.

Default

N/A.

Usage Guidelines

N/A.

Example

The following command removes the internal local clock from the clock source list:

```
configure ntp local-clock none
```

History

This command was first available in ExtremeXOS 12.7.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ntp local-clock stratum

```
configure ntp local-clock stratum stratum_number
```

Description

Configures the internal local clock with a stratum number. The stratum number defines the distance from the reference clock. The lower the number, the closer the switch is to the reference clock.

Syntax Description

<i>stratum_number</i>	Specifies the distance from the reference clock from 2 through 16, with 2 being closest and 16 being the farthest away.
-----------------------	---

Default

The local clock is disabled by default.

Usage Guidelines

The internal local clock is configured as a clock source with a given stratum number. Because the local clock is not as reliable as an external clock source with GPS or CDMA, the stratum number should be higher than the stratum number of the external clock source to allow the system to acquire the most reliable clock information from the clock source lists.

Example

The following command configures the local clock with a stratum number of 3:

```
configure ntp local-clock stratum 3
```

History

This command was first available in ExtremeXOS 12.7.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ntp restrict-list

```
configure ntp restrict-list [add | delete] network {mask} [permit |  
deny] [{vr} vr_name]
```

Description

Restricts a host or block of client IP addresses from getting NTP service. When NTP is enabled over a VLAN, an NTP server is configured, or a broadcast NTP server is in a VLAN, the VLAN's IP block or NTP server's IP address is automatically added into the system with a permit action.

Syntax Description

add	Restricts a client from getting NTP service.
delete	Removes a client from the restrict list.
<i>network</i>	Specifies a host or block of IP addresses.
<i>mask</i>	Specifies the subnet mask of the network.
permit	Specifies that a particular block of client IP addresses is permitted to get NTP service.
deny	Specifies that a particular block of client IP addresses is denied NTP service.
vr	Specifies VRs for NTP service.
<i>vr_name</i>	Specifies the VR name for allowing/denying NTP service. If no VR name is specified, the current command context is used.

Default

All addresses are denied by default.

If no VR name is specified, the current command context is used.

Usage Guidelines

N/A.

Example

The following command restricts a block of client IP addresses from getting NTP service:

```
configure ntp restrict-list add 132.25.82.3 deny
```

History

This command was first available in ExtremeXOS 12.7.

The **vr** keyword was added in ExtremeXOS 22.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ntp server/peer add

```
configure ntp [server | peer] add [ip_address | host_name] {key keyid}  
    {option [burst | initial-burst]} [{vr} vr_name]
```

Description

Configures an NTP server or peer.

Syntax Description

<i>ip_address</i>	Specifies the IP address of the NTP server or peer.
<i>host_name</i>	Specifies the host name of the NTP server or peer.
<i>keyid</i>	Specifies the key ID as a value from 1 to 65534.
burst	Follows the same burst mechanism when an NTP server is reachable.
initial-burst	Allows the system to send six burst packets when an NTP server becomes unreachable (discovered but unreachable).
vr	Specifies VR.
<i>vr_name</i>	Specifies the VR name. If no VR name is specified, the current command context is used.

Default

If no VR name is specified, the current command context is used.

Usage Guidelines

The initial-burst option is useful when a fast time synchronization is required at the initial stage.

Example

The following command adds an NTP server named “Missouri” with key 5 and an initial burst:

```
configure ntp server add Missouri key 5 initial-burst
```

History

This command was first available in ExtremeXOS 12.7.

The **vr** keyword was added in ExtremeXOS 22.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ntp server/peer delete

```
configure ntp [server | peer] delete [ip_address | host_name]
```

Description

Removes an NTP server or peer from external clock source lists.

Syntax Description

<i>ip_address</i>	Specifies the IP address of the NTP server or peer.
<i>host_name</i>	Specifies the host name of the NTP server or peer.

Default

N/A.

Usage Guidelines

N/A.

Example

The following command removes an NTP peer Missouri from external clock source lists

```
configure ntp peer delete Missouri
```

History

This command was first available in ExtremeXOS 12.7.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ospf bfd

```
configure ospf vlan vlan-name bfd on | off
```

Description

Configures BFD for OSPFv2.

Syntax Description

bfd	Bidirectional forwarding detection.
on	Turn on BFD for <i>OSPF</i> interface.
off	Turn off BFD for OSPF interface.

Default

Off.

Usage Guidelines

Use this command to turn BFD protection on or off on a specific OSPF interface.

The following example configures BFD protection on for *VLAN* 1:

Example

```
configure ospf vlan1 bfd on
```

History

This command was first available in ExtremeXOS 15.3.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ospf add virtual-link

```
configure ospf add virtual-link router-identifier area-identifier
```

Description

Adds a virtual link connected to another ABR.

Syntax Description

<i>router-identifier</i>	Specifies the router ID of the other end of the link.
<i>area-identifier</i>	Specifies an <i>OSPF</i> 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:

- router-identifier—Far-end router interface number.
- area-identifier—Transit area used for connecting the two end-points. The transit area cannot have the area identifier 0.0.0.0. and cannot be a stub area or an NSSA.

Example

The following command configures a virtual link between the two interfaces:

```
configure ospf add virtual-link 10.1.2.1 10.1.0.0
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [*Switch Engine 32.3 Feature License Requirements*](#) document.

configure ospf add vlan area

```
configure ospf add vlan [vlan-name | all] area area-identifier {passive}
```

Description

Enables *OSPF* on one or all *VLANs* (router interfaces).

Syntax Description

<i>vlan-name</i>	Specifies a VLAN name.
all	Specifies all VLANs.
<i>area-identifier</i>	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

Not applicable.

Example

The following command enables OSPF on a VLAN named accounting:

```
configure ospf add vlan accounting area 0.0.0.1
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [*Switch Engine 32.3 Feature License Requirements*](#) document.

configure ospf add vlan area link-type

```
configure ospf add vlan vlan-name area area-identifier link-type [auto |  
    broadcast | point-to-point] {passive}
```

Description

Configures the [*OSPF*](#) link type.

Syntax Description

<i>vlan-name</i>	Specifies a <i>VLAN</i> name.
<i>area-identifier</i>	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:

```
configure ospf add vlan accounting area 0.0.0.1 link-type auto
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [*Switch Engine 32.3 Feature License Requirements*](#) document.

configure ospf area add range

```
configure ospf area area-identifier add range [ip-address ip-mask |  
ipNetmask] [advertise | noadvertise] [type-3 | type-7]
```

Description

Configures a range of IP addresses in an [*OSPF*](#) area to be aggregated.

Syntax Description

<i>area-identifier</i>	Specifies an OSPF area.
<i>ip-address</i>	Specifies an IP address
<i>ip-mask</i>	Specifies a subnet mask.
<i>ipNetmask</i>	Specifies IP address / Netmask.
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:

```
configure ospf area 1.2.3.4 add range 10.1.2.0/24 advertise type-3
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [***Switch Engine 32.3 Feature License Requirements***](#) document.

configure ospf area delete range

```
configure ospf area area-identifier delete range [ip-address ip-mask |  
ipNetmask]
```

Description

Deletes a range of aggregated IP addresses in an [*OSPF*](#) area.

Syntax Description

<i>area-identifier</i>	Specifies an OSPF area.
<i>ip-address</i>	Specifies an IP address.
<i>ip-mask</i>	Specifies a subnet mask.
<i>ipNetmask</i>	Specifies IP address / Netmask.

Default

N/A.

Usage Guidelines

Not applicable.

Example

The following command deletes an aggregated IP address range:

```
configure ospf area 1.2.3.4 delete range 10.1.2.0/24
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [*Switch Engine 32.3 Feature License Requirements*](#) document.

configure ospf area external-filter

```
configure ospf area area-identifier external-filter [policy-map | none]
```

Description

Configures an external filter policy.

Syntax Description

<i>area-identifier</i>	Specifies the <i>OSPF</i> target area.
<i>policy-map</i>	Specifies a policy.
none	Specifies not to apply an external filter (removes the existing policy, if any).

Default

N/A.

Usage Guidelines

For switches configured to support multiple OSPF areas (an ABR function), a policy can be applied to an OSPF area that filters a set of OSPF external routes from being advertised into that area.

Using the none mode specifies that no external filter is applied.

Example

The following command configures an external filter policy, nosales:

```
configure ospf area 1.2.3.4 external-filter nosales
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure ospf area interarea-filter

```
configure ospf area area-identifier interarea-filter [policy-map | none]
```

Description

Configures a global inter-area filter policy.

Syntax Description

<i>area-identifier</i>	Specifies the OSPF target area.
<i>policy-map</i>	Specifies a policy.
none	Specifies not to apply an interarea filter.

Default

N/A.

Usage Guidelines

For switches configured to support multiple OSPF areas (an ABR function), a policy 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, nosales:

```
configure ospf area 0.0.0.6 interarea-filter nosales
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure ospf area normal

```
configure ospf area area-identifier normal
```

Description

Configures an OSPF area as a normal area.

Syntax Description

<i>area-identifier</i>	Specifies an <i>OSPF</i> 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:

```
configure ospf area 10.1.0.0 normal
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [*Switch Engine 32.3 Feature License Requirements*](#) document.

configure ospf area nssa stub-default-cost

```
configure ospf area area-identifier nssa [summary | nosummary] stub-  
default-cost cost {translate}
```

Description

Configures an *OSPF* area as an NSSA.

Syntax Description

<i>area-identifier</i>	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.
<i>cost</i>	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 routes originating from an ASBR connected to the NSSA can be advertised within the NSSA.
- External routes originating from the NSSA can be propagated to other areas, including the backbone area, if translated to type 5 LSAs.

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:

```
configure ospf area 10.1.1.0 nssa summary stub-default-cost 10 translate
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [**Switch Engine 32.3 Feature License Requirements**](#) document.

configure ospf area stub stub-default-cost

```
configure ospf area area-identifier stub [summary | nosummary] stub-  
default-cost cost
```


Description

Configures an *OSPF* area as a stub area.

Syntax Description

<i>area-identifier</i>	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.
<i>cost</i>	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:

```
configure ospf area 0.0.0.6 stub nosummary stub-default-cost 10
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the *Switch Engine 32.3 Feature License Requirements* document.

configure ospf area timer

```
configure ospf area area-identifier timer retransmit-interval transit-delay hello-interval dead-interval {wait-timer-interval}
```

Description

Configures the timers for all interfaces in the same *OSPF* area.

Syntax Description

<i>area-identifier</i>	Specifies an OSPF area.
<i>retransmit-interval</i>	Specifies the length of time that the router waits before retransmitting an LSA that is not acknowledged. The range is 1–3,600 seconds.
<i>transit-delay</i>	Specifies the length of time it takes to transmit an LSA packet over the interface. The range is 1–3,600 seconds.
<i>hello-interval</i>	Specifies the interval at which routers send hello packets. The range is 1–65,535 seconds.
<i>dead-interval</i>	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. The range is 1–2,147,483,647 seconds.
<i>wait-timer-interval</i>	Specifies the interval between the interface coming up and the election of the DR and BDR. Usually equal to the dead timer interval.

Default

- *retransmit interval*—Default: 5
- *transit delay*—Default: 1
- *hello interval*—Default: 10
- *dead interval*—Default: 40
- *wait timer interval*—Default: dead interval

Usage Guidelines

Configuring OSPF timers 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.

Specify the following:

- *retransmit interval*—If you set an interval that is too short, unnecessary retransmissions will result.
- *transit delay*—The transit delay must be 1 second or greater.
- *hello interval*—Smaller times allow routers to discover each other more quickly, but also increase network traffic.
- *dead interval*—This interval should be a multiple of the hello interval.
- *wait timer interval*—This interval is required by the OSPF standard to be equal to the router dead interval. Under some circumstances, setting the wait interval 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 hello interval. The default value is equal to the router dead interval.

Example

The following command sets the timers in area 0.0.0.2:

```
# configure ospf area 0.0.0.2 timer 10 1 20 200
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [*Switch Engine 32.3 Feature License Requirements*](#) document.

configure ospf ase-limit

```
configure ospf ase-limit number {timeout seconds}
```

Description

Configures the AS-external LSA limit and overflow duration associated with [*OSPF*](#) database overflow handling.

Syntax Description

<i>number</i>	Specifies the number of external routes that can be held in a link-state database.
<i>seconds</i>	Specifies a duration for which the system has to remain in the overflow state.

Default

The default for timeout is 0, which indicates that once the router goes into overflow state, it stays there until OSPF is disabled and then re-enabled.

Usage Guidelines

Not applicable.

Example

The following command configures the AS-external LSA limit and overflow duration:

```
configure ospf ase-limit 50000 timeout 1800
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [*Switch Engine 32.3 Feature License Requirements*](#) document.

configure ospf ase-summary add

```
configure ospf ase-summary add [ipaddress ip-mask | ipNetmask] cost cost
                               {tag number}
```

Description

Aggregates AS-external routes in a specified address range.

Syntax Description

<i>ipaddress</i>	Specifies an IP address.
<i>ip-mask</i>	Specifies a subnet mask.
<i>ipNetmask</i>	Specifies IP address / Netmask.
<i>cost</i>	Specifies a metric that will be given to the summarized route.
tag	Specifies an <i>OSPF</i> external route tag.

Default

N/A.

Usage Guidelines

This command is only valid on an ASBR.

Example

The following command summarizes AS-external routes:

```
configure ospf ase-summary add 175.1.0.0/16 cost 10
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure ospf ase-summary delete

```
configure ospf ase-summary delete [ip-address ip-mask | ipNetmask]
```

Description

Deletes an aggregated [OSPF](#) external route.

Syntax Description

<i>ip-address</i>	Specifies an IP address.
<i>ip-mask</i>	Specifies a subnet mask.
<i>ipNetmask</i>	Specifies IP address / Netmask.

Default

N/A.

Usage Guidelines

This command is only valid on an ASBR.

Example

The following command deletes the aggregated AS-external route:

```
configure ospf ase-summary delete 175.1.0.0/16
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure ospf authentication

```
configure ospf [vlan [vlan-name | all ] | area area-identifier |  
               virtual-link router-identifier area-identifier] authentication [ none  
               | encrypted simple-password encrypted-simple-password | simple-
```

```
password { simple-password } | encrypted md5 md5-key-id encrypted-md5-key | md5 md5-key-id { md5-key }
```

Description

Specifies the authentication password (up to eight characters) or RSA Data Security, Inc. [MD5](#) Message-Digest Algorithm key for one or all interfaces in a specific area or a virtual link.

Syntax Description

<i>vlan-name</i>	Specifies a VLAN name.
all	Specifies all VLANs
<i>router-identifier</i>	Specifies the router ID of the remote router.
<i>area-identifier</i>	Specifies an OSPF area.
encrypted	Indicates that the password (or key) is already encrypted (do not use this option).
<i>simple-password</i>	Specifies an authentication password (up to 8 ASCII characters).
<i>md5_key_id</i>	Specifies a RSA Data Security, Inc. MD5 Message-Digest Algorithm key, from 0-255.
<i>md5_key</i>	Specifies a numeric value from 0-65,536. Can also be alphanumeric, up to 26 characters.
none	Disables authentication.

Default

N/A.

Usage Guidelines

The *md5_key* is a numeric value with the range 0 to 65,536 or alphanumeric. When the OSPF area is specified, authentication information is applied to all OSPF interfaces within the area.

The **encrypted** option is used by the switch when generating a configuration file and when parsing a switch-generated configuration file. Do not select the **encrypted** option in the CLI.

Example

The following command configures RSA Data Security, Inc. MD5 Message-Digest Algorithm authentication on the VLAN subnet_26:

```
configure ospf vlan subnet_26 authentication md5 32 test
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure ospf cost

```
configure ospf [area area-identifier | vlan [vlan-name | all]] cost
[automatic | cost]
```

Description

Configures the cost metric of one or all interface(s) or an area.

Syntax Description

<i>area-identifier</i>	Specifies an OSPF area.
<i>vlan-name</i>	Specifies a VLAN name.
all	Specifies all VLANs.
automatic	Determine the advertised cost from the OSPF metric table.
<i>cost</i>	Specifies the cost metric.

Default

The default cost is automatic.

Usage Guidelines

The range is 1 through 65535.

Example

The following command configures the cost metric of the VLAN accounting:

```
configure ospf vlan accounting cost 10
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure ospf delete virtual-link

```
configure ospf delete virtual-link router-identifier area-identifier
```

Description

Removes a virtual link.

Syntax Description

<i>router-identifier</i>	Specifies the router ID of the other end of the link.
<i>area-identifier</i>	Specifies an <i>OSPF</i> area.

Default

N/A.

Usage Guidelines

None.

Example

The following command deletes a virtual link:

```
configure ospf delete virtual-link 10.1.2.1 10.1.0.0
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the *Switch Engine 32.3 Feature License Requirements* document.

configure ospf delete vlan

```
configure ospf delete vlan [vlan-name | all]
```

Description

Disables *OSPF* on one or all VLANs (router interfaces).

Syntax Description

<i>vlan-name</i>	Specifies a VLAN name.
all	Specifies all VLANs.

Default

N/A.

Usage Guidelines

Not applicable.

Example

The following command disables OSPF on VLAN accounting:

```
configure ospf delete vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure ospf import-policy

```
configure ospf import-policy [policy-map | none]
```

Description

Configures the import policy for [OSPF](#).

Syntax Description

<i>policy-map</i>	Specifies the policy.
-------------------	-----------------------

Default

No policy.

Usage Guidelines

An import policy is used to modify route attributes while adding OSPF routes to the IP route table. This command provides the flexibility of using import policy to determine the routes to be added to or removed from the routing table. In order to prevent a route being added to the routing table, the policy file must contain a matching rule with action “deny”. If there is no matching rule for a particular route, or the keyword “deny” is missing in the rule, the default action is “permit”, which means that route will be installed into the routing table.

Use the **none** option to remove an import policy.

If a policy rule set the cost to be greater than 65535, OSPF limits the metric of any matching routes to be 65535.

Example

The following example applies the policy campuseast to OSPF routes:

```
configure ospf import-policy campuseast
```

History

This command was first available in ExtremeXOS 10.1.

Beginning in ExtremeXOS 15.7, this command allows Import Policy to be used by OSPFv2 to install routes selectively into the switch routing table.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [***Switch Engine 32.3 Feature License Requirements***](#) document.

NEW! configure ospf instanceid

```
configure ospf instanceid instance-identifier
```

Description

Configures the Instance Identifier for OSPF instances.

Syntax Description

<i>instance-identifier</i>	Specifies the Instance Identifier. Default is 0. Range is 0-255.
----------------------------	--

Default

Zero.

Usage Guidelines

RFC 6549 specifies an 8-bit field in the OSPF header by splitting the previous 16-bit Authentication Type field into Instance ID and AuType. If Instance ID is configured, then all OSPF packets sent from the VR will have the Instance ID field set in the packet headers. If Instance IDs in the received packets do not match the configuration of the VR, then the packets will be discarded.

Example

The following example configures an Instance Identifier of 100:

```
configure ospf instanceid 100
```

History

This command was first available in ExtremeXOS 32.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ospf lsa-batch-interval

```
configure ospf lsa-batch-interval seconds
```

Description

Configures the *OSPF* LSA batching interval.

Syntax Description

<i>seconds</i>	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 batch interval to a value of 100 seconds:

```
configure ospf lsa-batch-interval 100
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [*Switch Engine 32.3 Feature License Requirements*](#) document.

configure ospf metric-table

```
configure ospf metric-table 10M cost_10m 100M cost_100m 1G cost_1g {2.5G
    cost_2_5g} {5G cost_5g} {10G cost_10g} {25Gcost_25g} {40G cost_40g}
    {50G cost_50g}{100G cost_100g}
```

Description

Configures the automatic interface costs for 10 Mbps, 100 Mbps, and 1 Gbps interfaces, and optionally, the 2.5 Gbps, 5 Gbps, 10 Gbps, 25 Gbps, 40 Gbps, 50 Gbps, and 100 Gbps interfaces.

Syntax Description

<i>cost</i>	Specifies the interface cost for the indicated interfaces.
-------------	--

Default

- 10 Mbps—The default cost is 10.
- 100 Mbps—The default cost is 5.
- 1 Gbps—The default cost is 4.
- 2.5 Gbps—The default cost is 3.
- 5 Gbps—The default cost is 3.
- 10 Gbps—The default cost is 2.
- 25 Gbps—The default cost is 2.
- 40 Gbps—The default cost is 2.
- 50 Gbps—The default cost is 2.
- 100 Gbps—The default cost is 1.

Usage Guidelines

Not applicable.

Example

The following command configures the automatic interface costs for 10 Mbps, 100 Mbps, and 1 Gbps interfaces:

```
configure ospf metric-table 10m 20 100m 10 1g 2
```

History

This command was first available in ExtremeXOS 10.1.

The 40 Gbps parameter was added in ExtremeXOS 12.6.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [*Switch Engine 32.3 Feature License Requirements*](#) document.

configure ospf priority

```
configure ospf [area area-identifier | vlan [vlan-name | all]] priority
               priority
```

Description

Configures the priority used in the designated router and backup designated router election algorithm for one or all [*OSPF*](#) interface(s) or for all the interfaces within the area.

Syntax Description

<i>area-identifier</i>	Specifies an OSPF area.
<i>vlan-name</i>	Specifies a <i>VLAN</i> name.
all	Specifies all VLANs.
<i>priority</i>	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 all the interfaces in area 1.2.3.4 to not be selected as the designated router:

```
configure ospf area 1.2.3.4 priority 0
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [***Switch Engine 32.3 Feature License Requirements***](#) document.

configure ospf restart grace-period

```
configure ospf restart grace-period seconds
```

Description

Configures the grace period sent out in Grace-LSAs and used by a restarting router.

Syntax Description

<i>seconds</i>	Grace period, in seconds. The default value is 120 seconds. Range is 1 to 1800 seconds.
----------------	---

Default

The default is 120 seconds.

Usage Guidelines

This command configures the grace period sent out to helper neighbor routers and used by the restarting router. The value of the grace period must be greater than the dead interval, and less than the LSA refresh time.

Example

The following command configures a router to send LSAs with a 240 second grace period during graceful *OSPF* restarts:

```
configure ospf restart grace-period 240
```

History

This command was first available in ExtremeXOS 11.3.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure ospf restart

```
configure ospf restart [none | planned | unplanned | both]
```

Description

Configures the router as a graceful [OSPF](#) restart router.

Syntax Description

none	Do not act as a graceful OSPF restart router.
planned	Only act as a graceful OSPF restart router for planned restarts.
unplanned	Only act as a graceful OSPF restart router for unplanned restarts.
both	Act as a graceful OSPF restart router for both planned and unplanned restarts.

Default

The default is none.

Usage Guidelines

This command configures the router as a graceful OSPF router. When configured for planned restarts, it will advertise Grace-LSAs before restarting (for example, during an upgrade of the OSPF module). When configured for unplanned restarts, it will advertise Grace-LSAs after restarting but before sending any Hellos. When configured for both, the router will advertise restarting regardless of whether the restart was planned or unplanned.

Example

The following command configures a router to perform graceful OSPF restarts only for planned restarts:

```
configure ospf restart planned
```

History

This command was first available in ExtremeXOS 11.3.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure ospf restart-helper

```
configure ospf [vlan [all | vlan-name] | area area-identifier |  
  virtual-link router-identifier area-identifier] restart-helper [none  
  | planned | unplanned | both]
```

Description

Configures the router as a graceful *OSPF* restart helper router.

Syntax Description

all	Specifies all <i>VLANs</i> .
<i>vlan-name</i>	Specifies a VLAN name.
<i>area-identifier</i>	Specifies an OSPF area.
<i>router-identifier</i>	Specifies the router ID of the remote router of the virtual link.
none	Do not act as a graceful OSPF restart helper router.
planned	Only act as a graceful OSPF restart helper router for planned restarts.
unplanned	Only act as a graceful OSPF restart helper router for unplanned restarts.
both	Act as a graceful OSPF restart helper router for both planned and unplanned restarts.

Default

The router default is none.

Usage Guidelines

This command configures the router as a graceful OSPF restart helper router for a single or multiple routers. When the router is acting as a helper, it will continue to advertise the restarting router as if it was fully adjacent.

One OSPF interface may not help more than one restarting router. An OSPF interface may not enter helper mode when the router is performing a graceful restart. All the interfaces to a neighbor router must be configured as graceful restart helpers, or the router will not support graceful restart for its neighbor.

Example

The following command configures a router to be a graceful OSPF helper router for planned restarts for all routers in area 10.20.30.40:

```
configure ospf area 10.20.30.40 restart-helper planned
```


History

This command was first available in ExtremeXOS 11.3.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [*Switch Engine 32.3 Feature License Requirements*](#) document.

configure ospf routerid

```
configure ospf routerid [automatic | router-identifier]
```

Description

Configures the [*OSPF*](#) router ID. If automatic is specified, the switch uses the highest IP interface address as the OSPF router ID.

Syntax Description

automatic	Specifies to use automatic addressing.
<i>router-identifier</i>	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.

Example

The following command sets the router ID:

```
configure ospf routerid 10.1.6.1
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [***Switch Engine 32.3 Feature License Requirements***](#) document.

configure ospf spf-hold-time

```
configure ospf spf-hold-time seconds
```

Description

Configures the minimum number of seconds between Shortest Path First (SPF) recalculations.

Syntax Description

<i>seconds</i>	Specifies a time in seconds. The range is 0 to 300 seconds.
----------------	---

Default

3 seconds.

Usage Guidelines

Not applicable.

Example

The following command configures the minimum number of seconds between Shortest Path First (SPF) recalculations:

```
configure ospf spf-hold-time 6
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [***Switch Engine 32.3 Feature License Requirements***](#) document.

configure ospf virtual-link timer

```
configure ospf virtual-link router-identifier area-identifier timer  
retransmit-interval transit-delay hello-interval dead-interval
```

Description

Configures the timers for a virtual link.

Syntax Description

<i>router-identifier</i>	Specifies the router ID of the other end of the link.
<i>area-identifier</i>	Specifies an OSPF area.
<i>retransmit-interval</i>	Specifies the length of time that the router waits before retransmitting an LSA that is not acknowledged. The range is 1–3,600 seconds.
<i>transit-delay</i>	Specifies the length of time it takes to transmit an LSA packet over the interface. The range is 1–3,600 seconds.
<i>hello-interval</i>	Specifies the interval at which routers send hello packets. The range is 1–65,535 seconds.
<i>dead-interval</i>	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. 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
- *wait timer interval*—Default: dead interval

Usage Guidelines

Configuring OSPF timers 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 and remote router ID 6.6.6.6:

```
configure ospf virtual-link 6.6.6.6 0.0.0.2 timer 10 1 20 200
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure ospf vlan area

```
configure ospf vlan vlan-name area area-identifier
```

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

<i>vlan-name</i>	Specifies a VLAN name.
<i>area-identifier</i>	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:

```
configure ospf vlan accounting area 0.0.0.6
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [*Switch Engine 32.3 Feature License Requirements*](#) document.

configure ospf vlan neighbor add

```
configure ospf vlan vlan-name neighbor add ip-address
```

Description

Configures the IP address of a point-to-point neighbor.

Syntax Description

<i>vlan-name</i>	Specifies a <i>VLAN</i> name.
<i>ip-address</i>	Specifies an IP address.

Default

N/A.

Usage Guidelines

None.

Example

The following command configures the IP address of a point-to-point neighbor:

```
configure ospf vlan accounting neighbor add 10.0.0.1
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [*Switch Engine 32.3 Feature License Requirements*](#) document.

configure ospf vlan neighbor delete

```
configure ospf vlan vlan-name neighbor delete ip-address
```

Description

Deletes the IP address of a point-to-point neighbor.

Syntax Description

<i>vlan-name</i>	Specifies a <u>VLAN</u> name.
<i>ip-address</i>	Specifies an IP address.

Default

N/A.

Usage Guidelines

None.

Example

The following command deletes the IP address of a point-to-point neighbor:

```
configure ospf vlan accounting neighbor delete 10.0.0.1
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the **Switch Engine 32.3 Feature License Requirements** document.

configure ospf vlan timer

```
configure ospf vlan [vlan-name | all] timer retransmit-interval transit-delay hello-interval dead-interval {wait-timer-interval}
```

Description

Configures the OSPF wait interval for a VLAN or all VLANs.

Syntax Description

<i>vlan-name</i>	Specifies a VLAN name.
all	Specifies all VLANs.

<i>retransmit-interval</i>	Specifies the length of time that the router waits before retransmitting an LSA that is not acknowledged. The range is 1-3,600.
<i>transit-delay</i>	Specifies the length of time it takes to transmit an LSA packet over the interface. The range is 1-3,600 seconds.
<i>hello-interval</i>	Specifies the interval at which routers send hello packets. The range is 1-65,535 seconds.
<i>dead-interval</i>	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. The range is 1-2,147,483,647.
<i>wait-timer-interval</i>	Specifies the interval between the interface coming up and the election of the DR and BDR. Usually equal to the dead timer interval.

Default

- *retransmit interval*—5 seconds.
- *transit delay*—1 second.
- *hello interval*—10 seconds.
- *dead interval*—40 seconds.
- *wait timer interval*—dead interval.

Usage Guidelines

Specify the following:

- *retransmit interval*—If you set an interval that is too short, unnecessary retransmissions will result.
- *transit delay*—The transit delay must be greater than 0.
- *hello interval*—Smaller times allow routers to discover each other more quickly, but also increase network traffic.
- *dead interval*—This interval should be a multiple of the hello interval.
- *wait timer interval*—This interval is required by the OSPF standard to be equal to the router dead interval. Under some circumstances, setting the wait interval 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 hello interval. The default value is equal to the router dead interval.

Example

The following command configures the OSPF wait interval on the VLAN accounting:

```
configure ospf vlan accounting timer 10 15 20 60 60
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [*Switch Engine 32.3 Feature License Requirements*](#) document.

configure ospfv3 add interface all

```
configure ospfv3 add [vlan | tunnel] all {instance-id instanceId} area
    area_identifier {passive}
```

Description

Enables [*OSPFv3*](#) on all VLANs or all tunnels (router interfaces).

Syntax Description

all	Specifies all IPv6 configured VLANs or all IPv6 tunnels.
<i>instanceId</i>	Specifies the instance ID for these interfaces. Range is 0 to 255.
<i>area_identifier</i>	Specifies the area to which the interfaces are assigned.
passive	Specifies to stop sending and receiving hello packets on this interface.

Default

OSPFv3 is disabled on the interfaces.

The default instance ID is 0.

Usage Guidelines

This command is used to enable the OSPFv3 protocol on all IPv6 configured VLANs or all IPv6 tunnels. The instance ID is used to control the selection of other routers as neighbors. The router will become a neighbor only with routers that have the same instance ID.

To change the instance ID associated with an interface, you must first remove the interface from the OSPFv3 area and then add it back with a different instance ID.

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 enables OSPFv3 on all IPv6 tunnels:

```
configure ospfv3 add tunnel all area 0.0.0.1
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure ospfv3 add interface

```
configure ospfv3 add [vlan vlan_name | tunnel tunnel_name] {instance-id
instanceId} area area_identifier link-type [auto | broadcast | point-
to-point] {passive}
```

Syntax Description

Enables [OSPFv3](#) on an interface.

Syntax Description

<i>vlan_name</i>	Specifies an IPv6 configured VLAN .
<i>tunnel_name</i>	Specifies an IPv6 tunnel.
<i>instanceId</i>	Specifies the instance ID for this interfaces. Range is 0 to 255.
<i>area_identifier</i>	Specifies the area to which the VLAN is assigned.
auto	Specifies to automatically determine the OSPFv3 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 hello packets on this interface.

Default

The default link-type is Auto.

The default instance ID is 0.

Usage Guidelines

This command is used to enable the OSPFv3 protocol on an IPv6 configured VLAN or an IPv6 tunnel. The instance ID is used to control the selection of other routers as neighbors. The router will become a neighbor only with routers that have the same instance ID.

To change the instance ID associated with an interface, you must first remove the interface from the OSPFv3 area and then add it back with a different instance ID.

The passive parameter indicates that the router only synchronizes and listens, and does not originate or send any new information on the interface.

Enable IPv6 forwarding before enabling OSPFv3; otherwise, you will receive a warning message.

You cannot change the **link-type** value while OSPFv3 is enabled on the interface.

Example

The following example adds the VLAN accounting (enabling OSPFv3 on the interface), to the area 0.0.0.1 with an instance ID of 2:

```
configure ospfv3 add vlan accounting instance-id 2 area 0.0.0.1 link-type auto
```

History

This command was first available in ExtremeXOS 11.2.

The **broadcast** and **point-to-point** link-type keywords were supported in ExtremeXOS 15.7.1.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [*Switch Engine 32.3 Feature License Requirements*](#) document.

configure ospfv3 add virtual-link

```
configure ospfv3 add virtual-link {routerid} router_identifier {area}  
                  area_identifier
```

Description

Adds a virtual link connected to another ABR.

Syntax Description

<i>router_identifier</i>	Specifies the router ID of the other end of the link.
<i>area_identifier</i>	Specifies the transit area identifier, a four-byte, dotted decimal number.

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:

- *router_identifier*—Far-end router identifier, a four-byte, dotted decimal number.
- *area_identifier*—Transit area used for connecting the two end-points. The transit area cannot have the area identifier 0.0.0.0. and cannot be a stub area or an NSSA.

Example

The following command configures a virtual link with router ID 10.1.2.1 through the transit area 10.1.0.0:

```
configure ospfv3 add virtual-link 10.1.2.1 10.1.0.0
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [*Switch Engine 32.3 Feature License Requirements*](#) document.

configure ospfv3 area add range

```
configure ospfv3 area area_identifier add range ipv6netmask [advertise |  
noadvertise] [inter-prefix | nssa]
```

Description

Configures a range of IPv6 addresses in an [*OSPFv3*](#) area to be aggregated.

Syntax Description

<i>area_identifier</i>	Specifies an OSPFv3 area, a four-byte, dotted decimal number.
<i>ipv6netmask</i>	Specifies an IPv6 address / prefix length.
advertise	Specifies to advertise the aggregated range of IPv6 addresses.
noadvert	Specifies not to advertise the aggregated range of IPv6 addresses.
inter-prefix	Specifies aggregate, inter-area-prefix LSAs.
nssa	NSSA LSAs.

Default

No OSPFv3 inter-area-prefix LSAs are configured.

Usage Guidelines

If advertised, the aggregated IPv6 range is exported as a single LSA by the ABR.

Example

The following command is used to summarize a certain range of IPv6 addresses within an area and export them out as a single address to area 0.0.0.1:

```
configure ospfv3 area 0.0.0.1 add range 2aaa:456:3ffe::/64 advertise inter-prefix
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [***Switch Engine 32.3 Feature License Requirements***](#) document.

configure ospfv3 area cost

```
configure ospfv3 area area_identifier cost [automatic | cost]
```

Description

Configures the cost of sending a packet to all interfaces belonging to an area.

Syntax Description

<i>area_identifier</i>	Specifies an <i>OSPFv3</i> area, a four-byte, dotted decimal number.
automatic	Determine the advertised cost from the OSPFv3 metric table.
<i>cost</i>	Specifies the cost metric. Range is 1 to 65535.

Default

The default cost is automatic.

Usage Guidelines

Use this command to set the cost of the links belonging to area manually, if the default cost needs to be overwritten. The interface cost is advertised as the link cost in router-LSA.

Example

The following command configures the cost of area 0.0.0.1 to 10. All the links of this area will inherit the area's cost value of 10.

```
configure ospfv3 area 0.0.0.1 cost 10
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure ospfv3 area delete range

```
configure ospfv3 area area_identifier delete range ipv6netmask [inter-prefix | nssa]
```

Description

Removes a range of IPv6 addresses in an [OSPFv3](#) area to be aggregated.

Syntax Description

<i>area_identifier</i>	Specifies an OSPFv3 area, a four-byte, dotted decimal number.
<i>ipv6netmask</i>	Specifies an IPv6 address / prefix length.
inter-prefix	Inter-Area-Prefix LSAs.
nssa	NSSA LSAs.

Default

No OSPFv3 inter-area-prefix LSAs are configured.

Usage Guidelines

If you attempt to delete a range that was not configured, you receive an error message.

Example

The following command is used to delete a summary network from area 0.0.0.1:

```
configure ospfv3 area 0.0.0.1 delete range 2aaa:456:3ffe::/64
```

History

This command was first available in ExtremeXOS 11.2.

The **inter-prefix** and **nssa** keywords were added in ExtremeXOS 21.1.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure ospfv3 area external-filter

```
configure ospfv3 area area_identifier external-filter [policy_map | none]
```

Description

Configures an external filter policy.

Syntax Description

<i>area_identifier</i>	Specifies the OSPFv3 target area.
<i>policy_map</i>	Specifies a policy.
none	Specifies not to apply an external filter (removes the existing policy, if any).

Default

N/A.

Usage Guidelines

For switches configured to support multiple OSPFv3 areas (an ABR function), a policy can be applied to an OSPFv3 area that filters a set of OSPFv3 external routes from being advertised into that area, in other words, filtering some of the inbound AS-external-LSAs.

OSPFv3 routers that do not have enough memory to hold the entire AS-external-LSA should configure an external area filter to drop part of the external-LSAs. Configuring this policy will enable routers with limited resources to be put into an OSPFv3 network.

Using the none mode specifies that no external filter is applied.

Policy files for this command will only recognize the following policy attributes:

- Match attributes:
 - *nlri IPv6-address/mask-len*
- Action (set) attributes
 - permit
 - deny

Any other policy attribute will not be recognized and will be ignored.

The following is an example of an external filter policy file:

```
entry one {  
    if match any{  
        nlri 2001:db8:3e5c::/48;
```

```

        nlri 2001:db8:2146:2341::/64;
    } then {
        deny;
    }
}

```

Example

The following command configures an external filter policy, nosales for area 1.2.3.4:

```
configure ospfv3 area 1.2.3.4 external-filter nosales
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [*Switch Engine 32.3 Feature License Requirements*](#) document.

configure ospfv3 area interarea-filter

```
configure ospfv3 area area_identifier interarea-filter [policy_map |
none]
```

Description

Configures an inter-area filter policy.

Syntax Description

<i>area_identifier</i>	Specifies the <i>OSPFv3</i> target area.
<i>policy_map</i>	Specifies a policy.
none	Specifies not to apply an inter-area filter (removes the existing policy, if any).

Default

N/A.

Usage Guidelines

ExtremeXOS OSPFv3 can apply an inter-area policy to filter some inter-area-prefix-LSAs and inter-area-router-LSAs from other areas. This can reduce the size of link state database of routers belonging to the area.

Using the none mode specifies that no external filter is applied.

Policy files for this command will only recognize the following policy attributes:

- Match attributes:
 - `nlri IPv6-address/mask-len`
- Action (set) attributes:
 - `permit`
 - `deny`

Any other policy attribute will not be recognized and will be ignored.

The following is an example of an inter-area filter policy file:

```
entry one {
    if match any{
        nlri 2001:db8:3e5c::/48;
        nlri 2001:db8:2146:2341::/64;
    } then {
        deny;
    }
}
entry two {
    if match any{
        nlri 2001:db8:444::/48;
        nlri 2001:db8:541f:65bd::/64;
    } then {
        permit;
    }
}
```

Example

The following command configures an inter-area filter policy, `nosales` for area 1.2.3.4:

```
configure ospfv3 area 1.2.3.4 interarea-filter nosales
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [***Switch Engine 32.3 Feature License Requirements***](#) document.

configure ospfv3 area normal

```
configure ospfv3 area area_identifier normal
```

Description

Configures an [*OSPFv3*](#) area as a normal area.

Syntax Description

<i>area_identifier</i>	Specifies an OSPFv3 area, a four-byte, dotted decimal number.
------------------------	---

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 OSPFv3 area as a normal area:

```
configure ospfv3 area 10.1.0.0 normal
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [*Switch Engine 32.3 Feature License Requirements*](#) document.

configure ospfv3 area nssa

```
configure ospfv3 area area-identifier nssa [nosummary | summary] stub-  
default-cost cost {translate}
```

Description

NSSAs are similar to the [*OSPFv3*](#) stub area configuration option, but have the following two additional capabilities:

- External routes originating from an ASBR connected to the NSSA can be advertised within the NSSA.
- External routes originating within the NSSA can be propagated to other areas if translated to AS-external LSAs. When configuring an OSPFv3 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, one of the ABRs for that NSSA is elected to perform translation.

Syntax Description

<i>area-identifier</i>	Area identifier.
nosummary	Inter-Area-Prefix LSAs prohibited.
summary	Inter-Area-Prefix LSAs allowed.
<i>cost</i>	Route metric.
translate	Always translate NSSA LSAs to AS-external LSAs.

Default

None.

Usage Guidelines

This command must specify the cost of the default route advertised into the NSSA.

History

This command was first available in ExtremeXOS 21.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ospfv3 area priority

```
configure ospfv3 area area_identifier priority priority
```

Description

Configures the priority used in the designated router and backup designated router election algorithm for all the interfaces within the area.

Syntax Description

<i>area_identifier</i>	Specifies an <u>OSPFv3</u> area, a four-byte, dotted decimal number.
<i>priority</i>	Specifies a priority range. The range is 0 through 255.

Default

The default setting is 1.

Usage Guidelines

When two routers are attached to a network, both attempt to become the designated router. The one with the higher priority takes precedence. If there is a tie, the router with the higher router ID takes

precedence. 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 all the interfaces in area 1.2.3.4 to not be selected as the designated router:

```
# configure ospfv3 area 1.2.3.4 priority 0
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [***Switch Engine 32.3 Feature License Requirements***](#) document.

On switches with a Core or Premier license, the non-zero interface priority takes effect; on switches with and Advanced Edge or Base license, the default interface priority is 0.

configure ospfv3 area stub

```
configure ospfv3 area area_identifier stub [summary | nosummary] stub-  
default-cost cost
```

Description

Configures an [*OSPFv3*](#) area as a stub area.

Syntax Description

<i>area_identifier</i>	Specifies an OSPFv3 area, a four-byte, dotted decimal number.
summary	Specifies that inter-area LSAs can be propagated into the area.
nosummary	Specifies that inter-area LSAs cannot be propagated into the area.
<i>cost</i>	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 consumption requirements on OSPFv3 routers.

Example

The following command configures an OSPFv3 area as a stub area:

```
configure ospfv3 area 0.0.0.6 stub nosummary stub-default-cost 10
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [**Switch Engine 32.3 Feature License Requirements**](#) document.

configure ospfv3 area timer

```
configure ospfv3 area area_identifier timer {retransmit-interval}
    retransmit_interval {transit-delay} transit_delay {hello-interval}
    hello_interval {dead-interval} dead_interval
```

Description

Configures the timers for all interfaces in the same [*OSPFv3*](#) area.

Syntax Description

<i>area_identifier</i>	Specifies an OSPFv3 area, a four-byte, dotted decimal number.
<i>retransmit_interval</i>	Specifies the length of time that the router waits before retransmitting an LSA that is not acknowledged. The range is 1 to 1,800 seconds.
<i>transit_delay</i>	Specifies the length of time it takes to transmit an LSA packet over the interface. The range is 1 to 1,800 seconds.
<i>hello_interval</i>	Specifies the interval at which routers send hello packets. The range is 1 to 65,535 seconds.
<i>dead_interval</i>	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. The range is 1 to 65,535 seconds.

Default

- Retransmit interval—Default: 5 seconds
- Transit delay—Default: 1 second
- Hello interval—Default: 10 seconds
- Dead interval—Default: 40 seconds

Usage Guidelines

Configuring OSPFv3 timers on a per-area basis is a shorthand for applying the timers to each VLAN and tunnel in the area at the time of configuration. If you add more VLANs or tunnels to the area, you must configure the timers for them explicitly.

Specify the following:

- Retransmit interval—If you set an interval that is too short, unnecessary retransmissions will result.
- Transit delay—The transit delay must be greater than 0.
- Hello interval—Smaller times allow routers to discover each other more quickly, but also increase network traffic.
- Dead interval—This interval should be a multiple of the hello interval.

The value of the dead interval and the hello interval must be same for all OSPFv3 routers connected to a common link. The value of the dead interval and the hello interval are advertised by OSPFv3 in Hello packets. The shorter the hello interval, the earlier topological changes will be detected, but more routing traffic will ensue.

The retransmit interval must be greater than the expected round trip delay between any two routers on the attached network. The setting of this parameter must be conservative, or needless retransmission will result.



Note

The wait interval for the interface is not separately configurable. It is always equal to the dead interval.

Example

The following command sets the timers in area 0.0.0.2:

```
configure ospfv3 area 0.0.0.2 timer 10 1 20 200
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the **Switch Engine 32.3 Feature License Requirements** document.

configure ospfv3 authentication (Authentication Trailer)

```
configure ospfv3 [{vlan} vlan-name | {tunnel} tunnel-name]  
    authentication [keychain keychain-name | none]
```

Description

Configures Authentication Trailer with a manual key to provide authentication on OSPFv3 interfaces.

Syntax Description

ospfv3	Specifies <i>OSPFv3</i> interface.
vlan	Specifies OSPFv3 VLAN.
<i>vlan-name</i>	Specifies an IPv6 configured <i>VLAN</i> .
tunnel	Specifies Layer 3 tunnel.
<i>tunnel-name</i>	Specifies Layer 3 tunnel name.
authentication	Specifies interface authentication.
keychain	Specifies the set of authentication keys.
<i>keychain-name</i>	Specifies the keychain name.
none	Specifies no authentication (default).

Default

If not specified, no authentication is applied.

Usage Guidelines

Users can only add keychains that are already present in the system. To add a keychain, run the command `create keychain keychain_name` .

Example

The following example for VLAN "vlan1" applies authentication type Authentication Trailer:

```
# configure ospfv3 vlan1 authentication keychain ospfv3-keys1
```

History

This command was first available in ExtremeXOS 31.3.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the *Switch Engine 32.3 Feature License Requirements* document.

configure ospfv3 authentication (IPsec)

```
configure ospfv3 [{vlan} vlan-name | {tunnel} tunnel-name]  
    authentication [none | ipsec spi spi esp-auth-algorithm algorithm key  
    [key-string | encrypted encrypted-key-string]
```

Description

Configures Internet Protocol Security (IPsec) with a manual key to provide authentication on OSPFv3 interfaces.

Syntax Description

ospfv3	Specifies <i>OSPFv3</i> interface.
vlan	Specifies OSPFv3 VLAN.
<i>vlan-name</i>	Specifies an IPv6 configured <i>VLAN</i> .
tunnel	Specifies Layer 3 tunnel.
<i>tunnel-name</i>	Specifies an Layer 3 tunnel name.
authentication	Specifies interface authentication.
none	Specifies no authentication (default).
ipsec spi	Specifies the authentication type is IPsec Encapsulating Security Payload (ESP) with manual key.
<i>spi</i>	Specifies Security Parameter Index value. Range is 256-4294967295.
esp-auth-algorithm	Specifies the ESP Authentication algorithm.
<i>algorithm</i>	Specifies the authentication algorithm. Supported authentication algorithms are hmac-sha-1 and hmac-sha-256.
key	Specifies the authentication key.
<i>key-string</i>	Specifies the key string in clear text. Both the ASCII string and hexadecimal string are supported, and hexadecimal string must begin with "0x".
encrypted	Specifies that the key string is in encrypted format.
<i>encrypted-key-string</i>	Specifies the encrypted key string. The encrypted key string must be enclosed in double quotes.

Default

If not specified, no authentication is applied.

Usage Guidelines

When configuring IPsec with manual key on an OSPFv3 VLAN, the exact same IPsec parameters (SPI, algorithm and key-string) must be specified on all routers connected to that VLAN.

To configure OSPFv3 virtual link authentication, run the command **ospfv3 virtual-link {routerid} router-identifier {area} area-identifier authentication [none | ipsec spi spi esp-auth-algorithm algorithm key [key-string | encrypted encrypted-key-string]**.

Example

The following example for VLAN "v1" applies authentication type IPsec with SPI "551" and algorithm "hmac-sha-256" with key "mykey":

```
# configure ospfv3 vlan v1 authentication ipsec spi 551 esp-auth-algorithm hmac-sha-256
key mykey
```

History

This command was first available in ExtremeXOS 31.2.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [*Switch Engine 32.3 Feature License Requirements*](#) document.

configure ospfv3 bfd

```
configure ospfv3 vlan vlan-name bfd on | off
```

Description

Configures BFD for [*OSPFv3*](#).

Syntax Description

bfd	Bidirectional forwarding detection
on	Turn on BFD for OSPFv3 interface.
off	Turn off BFD for OSPFv3 interface.

Default

Off.

Usage Guidelines

Use this command to turn on or off BFD protection on a specific OSPFv3 interface.

The following example configures BFD protection on for [*VLAN*](#) 1:

Example

```
# configure ospfv3 vlan1 bfd on
```

History

This command was first available in ExtremeXOS 15.3.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ospfv3 delete interface

```
configure ospfv3 delete [vlan vlan_name | tunnel tunnel_name | [vlan |
tunnel] all]
```

Description

Disables *OSPFv3* on one or all *VLANs* or tunnels (router interfaces).

Syntax Description

<i>vlan_name</i>	Specifies an IPv6 configured VLAN.
<i>tunnel_name</i>	Specifies an IPv6 tunnel.
all	Specifies all VLANs, or tunnels.

Default

N/A.

Usage Guidelines

None.

Example

The following command disables OSPFv3 on VLAN accounting:

```
configure ospfv3 delete vlan accounting
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the *Switch Engine 32.3 Feature License Requirements* document.

configure ospfv3 delete virtual-link

```
configure ospfv3 delete virtual-link {routerid} router_identifier {area}
area_identifier
```

Description

Deletes a virtual link connected to another ABR.

Syntax Description

<i>router_identifier</i>	Specifies the router ID of the other end of the link.
<i>area_identifier</i>	Specifies the transit area identifier, a four-byte, dotted decimal number.

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:

- Router-identifier—Far-end router identifier, a four-byte, dotted decimal number.
- Area-identifier—Transit area used for connecting the two end-points. The transit area cannot have the area identifier 0.0.0.0. and cannot be a stub area or an NSSA.

Example

The following command deletes a virtual link with router ID 10.1.2.1 through the transit area 10.1.0.0:

```
configure ospfv3 delete virtual-link 10.1.2.1 10.1.0.0
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [**Switch Engine 32.3 Feature License Requirements**](#) document.

configure ospfv3 import-policy

```
configure ospfv3 import-policy [policy_map | none]
```

Description

Configures the import policy for OSPFv3.

Syntax Description

<code>policy_map</code>	Specifies the policy.
-------------------------	-----------------------

Default

No policy.

Usage Guidelines

An import policy is used to modify route attributes while adding OSPFv3 routes to the IPv6 route table. This command provides the flexibility of using import policy to determine the routes to be added to or removed from the routing table. In order to prevent a route being added to the routing table, the policy file must contain a matching rule with action “deny”. If there is no matching rule for a particular route, or the keyword “deny” is missing in the rule, the default action is “permit”, which means that route will be installed into the routing table.

Use the **none** option to remove the policy association.

Policy files for this command will recognize only the following policy attributes:

- Match attributes:
 - `nlri IPv6-address/mask-len`
 - `route-origin [ospf | ospf-extern1 | ospf-extern2 | ospf-inter | ospf-intra]`
- Action (set) attributes
 - `cost cost`
 - `tag number`
 - `deny`

Any other policy attribute will not be recognized and will be ignored.

Example

The following example applies the policy `campuseast` to OSPFv3 routes:

```
configure ospfv3 import-policy campuseast
```

History

This command was first available in ExtremeXOS 11.2.

Beginning in ExtremeXOS 15.7, this command allows Import Policy to be used by OSPFv3 to install routes selectively into the switch routing table.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [**Switch Engine 32.3 Feature License Requirements**](#) document.

configure ospfv3 interface area

```
configure ospfv3 [vlan vlan_name | tunnel tunnel_name] area
area_identifier
```

Description

Moves an interface from one *OSPFv3* area to another.

Syntax Description

<i>vlan_name</i>	Specifies an IPv6 configured <i>VLAN</i> .
<i>tunnel_name</i>	Specifies an IPv6 tunnel.
<i>area_identifier</i>	Specifies an OSPFv3 area, a four-byte, dotted decimal number.

Default

N/A.

Usage Guidelines

Use this command to move an already configured interface from one area to another. The instance ID associated with the interface will be unchanged.

Example

The following command moves the VLAN accounting to the OSPFv3 area 0.0.0.6:

```
configure ospfv3 vlan accounting area 0.0.0.6
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the *Switch Engine 32.3 Feature License Requirements* document.

configure ospfv3 interface cost

```
configure ospfv3 [vlan vlan_name | tunnel tunnel_name | [vlan | tunnel]
all]] cost [automatic | cost]
```

Description

Configures the cost of one or all interface(s).

Syntax Description

<i>vlan_name</i>	Specifies an IPv6 configured VLAN .
<i>tunnel_name</i>	Specifies an IPv6 tunnel.
all	Specifies all IPv6 configured VLANs or all IPv6 tunnels.
automatic	Determine the advertised cost from the OSPFv3 metric table.
<i>cost</i>	Specifies the cost metric. Range is 1 to 65535.

Default

The default cost is automatic.

Usage Guidelines

Use this command to set the cost of an interface (a VLAN or tunnel) manually, if the default cost needs to be overwritten. The interface cost is advertised as the link cost in router-LSA.

Example

The following command configures the cost metric of the VLAN accounting:

```
configure ospfv3 vlan accounting cost 10
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure ospfv3 interface priority

```
configure ospfv3 [vlan vlan_name | tunnel tunnel_name | [vlan | tunnel]  
                 all] priority priority
```

Description

Configures the priority used in the designated router and backup designated router election algorithm for one or all [OSPFv3](#) interface(s).

Syntax Description

<i>vlan_name</i>	Specifies an IPv6 configured VLAN .
<i>tunnel_name</i>	Specifies an IPv6 tunnel.
all	Specifies all IPv6 configured VLANs or all IPv6 tunnels.
<i>priority</i>	Specifies a priority range. The range is 0 through 255.

Default

The default setting is 1.

Usage Guidelines

When two routers are attached to a network, both attempt to become the designated router. The one with the higher priority takes precedence. If there is a tie, the router with the higher router ID takes precedence. 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 priority of the interface VLAN corporate to 10:

```
# configure ospfv3 vlan corporate priority 10
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

On switches with a Core or Premier license, the non-zero interface priority takes effect; on switches with an Advanced Edge or Base license, the default interface priority is 0.

configure ospfv3 interface timer

```
configure ospfv3 [vlan vlan_name | tunnel tunnel_name | [vlan  
| tunnel] all] timer {retransmit-interval} retransmit_interval  
{transit-delay} transit_delay {hello-interval} hello_interval {dead-  
interval} dead_interval
```

Description

Configures the timers for all interfaces in the same [OSPFv3](#) area.

Syntax Description

<i>vlan_name</i>	Specifies an IPv6 configured <u>VLAN</u> .
<i>tunnel_name</i>	Specifies an IPv6 tunnel.
all	Specifies all IPv6 configured VLANs or all IPv6 tunnels.
<i>retransmit_interval</i>	Specifies the length of time that the router waits before retransmitting an LSA that is not acknowledged. The range is 1 to 3600 seconds.
<i>transit_delay</i>	Specifies the length of time it takes to transmit an LSA packet over the interface. The range is 1 to 3600 seconds.
<i>hello_interval</i>	Specifies the interval at which routers send hello packets. The range is 1 to 65535 seconds.
<i>dead_interval</i>	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. The range is 1 to 65535 seconds.

Default

- Retransmit interval—Default: 5 seconds.
- Transit delay—Default: 1 second.
- Hello interval—Default: 10 seconds.
- Dead interval—Default: 40 seconds.

Usage Guidelines

Use this command to configure the OSPFv3 timers on a per-interface basis.

Specify the following:

- retransmit interval—If you set an interval that is too short, unnecessary retransmissions will result.
- transit delay—The transit delay must be greater than 0.
- hello interval—Smaller times allow routers to discover each other more quickly, but also increase network traffic.
- dead interval—This interval should be a multiple of the hello interval.

The value of the dead interval and the hello interval must be same for all OSPFv3 routers connected to a common link. The value of the dead interval and the hello interval are advertised by OSPFv3 in Hello packets. The shorter the hello interval, the earlier topological changes will be detected, but more routing traffic will ensue.

The retransmit interval must be greater than the expected round trip delay between any two routers on the attached network. The setting of this parameter must be conservative, or needless retransmission will result.



Note

The wait interval for the interface is not separately configurable. It is always equal to the dead interval.

Example

The following command sets the timers for the VLAN corporate:

```
configure ospfv3 vlan corporate timer retransmit-interval 10 transit-delay 2 hello-
interval 20 dead-interval 80
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [***Switch Engine 32.3 Feature License Requirements***](#) document.

configure ospfv3 lsa-batch-interval

```
configure ospfv3 lsa-batch-interval seconds
```

Description

This command configures the LSA batch interval. LSAs added during this interval are batched together for update.

Syntax Description

<i>seconds</i>	Interval in seconds. Range is 0 to 600. (Default 0, not batched).
----------------	---

Default

0.

Usage Guidelines

The range is 0 to 600 seconds.

Example

The following example shows the output of the `show ospfv3` command including the LSA batch interval output:

```
# show ospfv3
OSPFv3           : Disabled           RouterId         : 0.0.0.0
RouterId Selection : Automatic         ASBR             : No
ABR              : No                  ExtLSAs          : 0
ExtLSAChecksum    : 0x0                OriginateNewLSAs : 0
ReceivedNewLSAs   : 0                  SpfHoldTime      : 3s
Num of Areas      : 1                  LSA Batch Interval : 0s
10M Cost          : 100                 100M Cost        : 50
1000M Cost (1G)   : 40                 2500M Cost (2.5G) : 40
```



```

5000M Cost (5G)      : 40
25000M Cost (25G)   : 20
50000M Cost (50G)   : 20
Graceful Restart    : None
Import Policy File   : none
SNMP Traps           : Disabled
Redistribute:
  Protocol           Status    Cost    Type    Tag     Policy
  direct              Disabled  20      2       ---    none
  e-bgp                Disabled  20      2       ---    none
  i-bgp                Disabled  20      2       ---    none
  ripng               Disabled  20      2       ---    none
  static              Disabled  20      2       ---    none
  isis-level-1        Disabled  20      2       ---    none
  isis-level-2        Disabled  20      2       ---    none
  isis-level-1-external Disabled  20      2       ---    none
  isis-level-2-external Disabled  20      2       ---    none
  host-mobility        Disabled  20      2       ---    none

```

History

This command was first available in ExtremeXOS 21.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ospfv3 metric-table

```

configure ospfv3 metric-table [{10M cost_10m} {100M cost_100m } {1G
cost_1g}{2.5G cost_2_5g} {5G cost_5g} {10G cost_10g} {25G cost_25g}
{40G cost_40g} {50G cost_50g} {100G cost_100g} ]

```

Description

Configures the optional interface costs for 10 Mbps, 100 Mbps, 1 Gbps, 2.5 Gbps, 5 Gbps, 10 Gbps, 25 Gbps, 40 Gbps, 50 Gbps, and 100 Gbps interfaces.

Syntax Description

<code>cost_x</code>	Specifies the interface cost for the indicated interfaces. Range is 1 to 65535.
---------------------	---

Default

- 10 Mbps—The default cost is 100.
- 100 Mbps—The default cost is 50.
- 1 Gbps—The default cost is 40.
- 2.5 Gbps—The default cost is 40.
- 5 Gbps—The default cost is 40.
- 10 Gbps—The default cost is 20.

- 25 Gbps—The default cost is 20.
- 40 Gbps—The default cost is 20.
- 50 Gbps—The default cost is 20.
- 100 Gbps—The default cost is 10.

Usage Guidelines

The value of the costs cannot be greater for higher speed interfaces. In other words, the following condition must be true:

cost_10m >= cost_100m >= cost_1g >= cost_2.5g >= cost_5g cost_10g >= cost_25g >= cost_40g >= cost_50g >= cost_100g

Example

The following command configures the automatic interface costs for 10 Mbps, 100 Mbps, 1 Gbps, 2.5 Gbps, 5 Gbps, 10 Gbps, 25 Gbps, 40 Gbps, 50 Gbps, and 100 Gbps interfaces:

```
configure ospfv3 metric-table 10M 110 100M 70 1G 50 2.5G 45 5G 40 10G 35 25G 30 40G 25
50G 20 100G 15
```

The following example displays the output of the show ospfv3 command:

```
# show ospfv3
OSPFv3                : Disabled          RouterId           : 0.0.0.0
RouterId Selection    : Automatic         ASBR               : No
ABR                   : No                ExtLSAs            : 0
ExtLSAChecksum        : 0x0              OriginateNewLSAs   : 0
ReceivedNewLSAs       : 0                SpfHoldTime        : 3s
Num of Areas           : 1                LSA Batch Interval : 0s
10M Cost               : 110              100M Cost          : 70
1000M Cost (1G)        : 50                2500M Cost (2.5G)  : 45
5000M Cost (5G)        : 40                10000M Cost (10G)  : 35
25000M Cost (25G)      : 30                40000M Cost (40G)  : 25
50000M Cost (50G)      : 20                100000M Cost (100G) : 15
Graceful Restart       : None              Grace Period        : 120s
Import Policy File     : none
SNMP Traps             : Disabled
Redistribute:
  Protocol      Status   Cost   Type  Tag   Policy
  direct        Disabled 20      2    ---   none
  e-bgp          Disabled 20      2    ---   none
  i-bgp          Disabled 20      2    ---   none
  ripng          Disabled 20      2    ---   none
  static         Disabled 20      2    ---   none
  isis-level-1   Disabled 20      2    ---   none
  isis-level-2   Disabled 20      2    ---   none
  isis-level-1-external Disabled 20      2    ---   none
  isis-level-2-external Disabled 20      2    ---   none
  host-mobility  Disabled 20      2    ---   none
```

History

This command was first available in ExtremeXOS 11.2.

The 40 Gbps parameter was added in ExtremeXOS 12.6.

The 2.5G, 5G, 25G and 50G speeds were added in ExtremeXOS 22.1.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [*Switch Engine 32.3 Feature License Requirements*](#) document.

configure ospfv3 restart

```
configure ospfv3 restart [none | planned | unplanned | both]
```

Description

This command configures the graceful restart behavior the router.

Syntax Description

none	Disable graceful restart.
planned	Support planned restart only.
unplanned	Support unplanned restart only.
both	Support both planned and unplanned restart.

Default

Graceful restart is disabled by default.

Usage Guidelines

When configured for planned restarts, it will support planned restarts (like process restart) and advertise Grace LSAs before restarting. When configured for unplanned restarts, it will support unplanned restarts (like failover in a stack) and advertise Grace LSAs after restarting but before sending any Hellos. When configured for both, the router will support both planned and unplanned restarts. Unplanned restarts and BFD configuration on interfaces are incompatible in ExtremeXOS. If both are enabled, an unplanned restart will fail.

History

This command was first available in ExtremeXOS 21.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ospfv3 restart grace-period

```
configure ospfv3 restart grace-period seconds
```

Description

This command configures the grace period sent out in Grace LSAs and used by a restarting router.

Syntax Description

<i>seconds</i>	Interval in seconds. Range is 1 to 1800.
----------------	--

Default

The default grace period is 120 seconds.

Usage Guidelines

The range is 1 to 1800 seconds. The grace period should be greater than hello interval and router dead interval of the OSPFv3 interfaces on the router.

History

This command was first available in ExtremeXOS 21.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ospfv3 restart-helper

```
configure ospfv3 [[vlan | tunnel] all | {vlan} vlan-name | {tunnel}  
  tunnel-name | area area-identifier] restart-helper [none | planned |  
  unplanned | both]
```

Description

This command configures graceful restart helper mode behavior of OSPFv3 interfaces for its neighbors. When an interface is acting as a helper, it will continue to advertise the restarting router as if it was fully adjacent.

Syntax Description

vlan	
all	Variable description, available options, and notes.
<i>vlan-name</i>	<u>VLAN</u> name.

area	OSPFv3 area.
<i>area-identifier</i>	Area identifier.
none	Disable helper mode.
planned	Support planned restart only.
unplanned	Support unplanned restart only.
both	Support both planned and unplanned restart.

Default

Restart helper mode is disabled by default.

Usage Guidelines

When the area option is used the command applies to all interfaces in the area at that time. One OSPFv3 interface may not help more than one restarting router at a time. An OSPFv3 interface may not enter helper mode when the router is performing a graceful restart. All the interfaces to a neighbor router must be configured as graceful restart helpers, or the router will not support graceful restart for its neighbor.

Restart Helper mode is displayed in the `show ospfv3 interfaces detail` output.

```
# show ospfv3 interfaces detail
Interface           : v100           Enabled           : ENABLED
Router              : ENABLED        AreaID            : 0.0.0.0
RouterID            : 10.1.1.2       Link Type         : point-to-point
Passive             : No             Cost              : 40/A
Priority             : 1              Transit Delay     : 1s
Hello Interval      : 10s            Rtr Dead Time    : 40s
Retransmit Interval : 5s            Wait Timer       : 40s
Interface ID        : 19             Instance ID      : 0
State               : P2P            Number of state chg : 1
Hello due in        : 7s            Number of events  : 2
Total Num of Nbrs   : 1              Nbrs in FULL State : 1
Hellos Rxed         : 127733         Hellos Txed       : 127739
DB Description Rxed  : 4              DB Description Txed : 3
LSA Request Rxed    : 1              LSA Request Txed  : 1
LSA Update Rxed     : 2121           LSA Update Txed   : 6156
LSA Ack Rxed        : 5962           LSA Ack Txed      : 2121
In Discards         : 0
DR RtId             : 0.0.0.0        BDR RtId         : 0.0.0.0
Restart Helper       : Both
Restart Helper Strict LSA Checking: Enabled
BFD Protection      : Off

Neighbors:
  RtrId: 10.1.1.1 IpAddr: fe80::204:96ff:fe51:ea8e Pri: 1 Type: Auto
  State: FULL DR: 0.0.0.0 BDR: 0.0.0.0 Dead Time: 00:00:31
  Options: 0x13 (-|R|-|E|V6) Opaque LSA: No
  Restart Helper Status: Off
  Last Restart Helper Exit Reason: None
  BFD Session State: None
```

History

This command was first available in ExtremeXOS 21.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ospfv3 routerid

```
configure ospfv3 routerid [automatic | router_identifier]
```

Description

Configures the *OSPFv3* router ID. If automatic is specified, the switch uses the highest IPv4 interface address as the OSPFv3 router ID.

Syntax Description

automatic	Specifies to use automatic addressing.
<i>router_identifier</i>	Specifies a router identifier, a four-byte, dotted decimal number.

Default

Automatic.

Usage Guidelines

Each switch that is configured to run OSPFv3 must have a unique router ID. The router ID is a four-byte, dotted decimal number, like an IPv4 address. Even though the IP address format has changed from IPv4 to IPv6, the router ID format has not. It is recommended that you manually set the router ID of the switches participating in OSPFv3, instead of having the switch automatically choose its router ID based on the highest interface IPv4 address (if it exists). Not performing this configuration in larger, dynamic environments could result in an older link-state database remaining in use.

This command is accepted only when OSPFv3 is globally disabled.



Note

Do not set the router ID to 0.0.0.0.

Example

The following command sets the router ID to 10.1.6.1:

```
configure ospfv3 routerid 10.1.6.1
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [*Switch Engine 32.3 Feature License Requirements*](#) document.

configure ospfv3 spf-hold-time

```
configure ospfv3 spf-hold-time seconds
```

Description

Configures the minimum number of seconds between Shortest Path First (SPF) recalculations.

Syntax Description

spf-hold-time	SPF hold time.
<i>seconds</i>	Specifies a time in seconds. The range is 0 to 300 seconds.

Default

3 seconds.

Usage Guidelines

Setting the interval too high will force [*OSPFv3*](#) to run SPF calculations less frequently. This will reduce the CPU load, but will cause delay in routes getting updated in the IPv6 routing table. Setting the interval too low will decrease the interval between SPF calculations, but will increase the processing load on CPU.

Example

The following command configures the minimum number of seconds between Shortest Path First (SPF) recalculations:

```
configure ospfv3 spf-hold-time 6
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure ospfv3 virtual-link authentication (Authentication Trailer)

```
configure ospfv3 virtual-link {routerid} router-identifier {area} area-
  identifier authentication [keychain keychain_name | none]
```

Description

Configure Authentication Trailer with a manual key to provide authentication on OSPFv3 virtual-links.

Syntax Description

ospfv3	Specifies OSPFv3 virtual-link.
virtual-link	OSPFv3 virtual link.
routerid	OSPFv3 router ID.
<i>router-identifier</i>	Specifies the router identifier of the advertising router.
area	OSPFv3 area.
<i>area-identifier</i>	Specifies an OSPFv3 area, a four-byte, dotted decimal number.
authentication	Specifies interface authentication.
keychain	Specifies the set of authentication keys.
<i>keychain-name</i>	Specifies the keychain name.
none	Specifies no authentication (default).

Default

If not specified, no authentication is applied.

Usage Guidelines

Users can only add keychains that are already present on the system. Keychains can be created using this command `create keychain keychain_name` .

Example

The following example for virtual-link with a router id of 10.1.1.3 and an area identifier of 1.1.1.1:

```
# configure ospfv3 virtual-link 10.1.1.3 area 1.1.1.1 authentication keychain ospfv3-
keys1
```

History

This command was first available in ExtremeXOS 31.3.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure ospfv3 virtual-link authentication

```
configure ospfv3 virtual-link {routerid} router-identifier {area} area-
  identifier authentication [none | keychain keychain-name | ipsec
  spi spi esp-auth-algorithm algorithm key [key-string | encrypted
  encrypted-key-string]
```

Description

Configure Internet Protocol Security (IPsec) with a manual key to provide authentication on OSPFv3 virtual-links.

Syntax Description

ospfv3	Specifies OSPFv3 virtual-link.
virtual-link	OSPFv3 virtual link.
routerid	OSPFv3 router ID.
<i>router-identifier</i>	Specifies the router identifier of the advertising router.
area	OSPFv3 area.
<i>area-identifier</i>	Specifies an OSPFv3 area, a four-byte, dotted decimal number.
authentication	Specifies interface authentication.
none	Specifies no authentication (default).
keychain	Specifies the authentication method is keychain.
<i>keychain-name</i>	Specifies the keychain name.
ipsec spi	Specifies the authentication type is IPsec Encapsulating Security Payload (ESP) with manual key.
<i>spi</i>	Specifies Security Parameter Index value. Range is 256-4294967295.
esp-auth-algorithm	Specifies the ESP Authentication algorithm.
<i>algorithm</i>	Specifies the authentication algorithm. Supported authentication algorithms are hmac-sha-1 and hmac-sha-256.
key	Specifies the authentication key,
<i>key-string</i>	Specifies the key string in clear text. Both the ASCII string and hexadecimal string are supported, and hexadecimal string must begin with "0x".
encrypted	Specifies that the key string is in encrypted format.
<i>encrypted-key-string</i>	Specifies the encrypted key string. The encrypted key string must be enclosed in double quotes.

Default

If not specified, no authentication is applied.

Usage Guidelines

When configuring IPsec with a manual key on an OSPFv3 virtual link, the exact same IPsec parameters (SPI, algorithm and key-string) must be specified on all routers connected to both sides of the virtual link.

To configure OSPFv3 VLAN authentication, run the command `configure ospfv3 [{vlan} vlan-name | {tunnel} tunnel-name] authentication [none | ipsec spi spi esp-auth-algorithm algorithm key [key-string | encrypted encrypted-key-string]]`.

Example

The following example for virtual-link "5.5.5.5 0.0.0.2" applies authentication type IPsec with SPI "1001" and algorithm "hmac-sha-1" with key "mykey":

```
# configure ospfv3 virtual-link 5.5.5.5 0.0.0.2 authentication ipsec spi 1001 esp-auth-
algorithm hmac-sha-1 key mykey
```

History

This command was first available in ExtremeXOS 31.2.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure ospfv3 virtual-link restart-helper

```
configure ospfv3 virtual-link {routerid} router-identifier {area} area-
identifier restart-helper [none | planned | unplanned | both]
```

Description

This command configures graceful restart helper mode behavior of [OSPFv3](#) interfaces for its neighbors. When an interface is acting as a helper, it continues to advertise the restarting router as if it was fully adjacent.

Syntax Description

virtual-link	OSPFv3 virtual link.
routerid	OSPFv3 router ID.
<i>router-identifier</i>	Router ID of neighbor OSPFv3 router.

area	OSPFv3 area.
<i>area-identifier</i>	Transit area ID of virtual link.
restart-helper	Graceful restart helper mode.
none	Disable helper mode (default).
planned	Support planned restart only.
unplanned	Support unplanned restart only.
both	Support both planned and unplanned restart.

Default

Helper mode is disabled by default.

Usage Guidelines

When the area option is used, the command applies to all interfaces in the area at that time. One OSPFv3 interface may not help more than one restarting router at a time. An OSPFv3 interface may not enter helper mode when the router is performing a graceful restart. All the interfaces to a neighbor router must be configured as graceful restart helpers, or the router does not support graceful restart for its neighbor.

Restart helper mode appears in the `show ospfv3 interfaces detail` output.

```
# show ospfv3 interfaces detail
Interface : v100 Enabled : ENABLED
Router : ENABLED AreaID : 0.0.0.0
RouterID : 10.1.1.2 Link Type : point-to-point
Passive : No Cost : 40/A
Priority : 1 Transit Delay : 1s
Hello Interval : 10s Rtr Dead Time : 40s
Retransmit Interval : 5s Wait Timer : 40s
Interface ID : 19 Instance ID : 0
State : P2P Number of state chg : 1
Hello due in : 7s Number of events : 2
Total Num of Nbrs : 1 Nbrs in FULL State : 1
Hellos Rxed : 127733 Hellos Txed : 127739
DB Description Rxed : 4 DB Description Txed : 3
LSA Request Rxed : 1 LSA Request Txed : 1
LSA Update Rxed : 2121 LSA Update Txed : 6156
LSA Ack Rxed : 5962 LSA Ack Txed : 2121
In Discards : 0
DR RtId : 0.0.0.0 BDR RtId : 0.0.0.0
Restart Helper : Both
Restart Helper Strict LSA Checking: Enabled
BFD Protection : Off

Neighbors:
RtrId: 10.1.1.1 IpAddr: fe80::204:96ff:fe51:ea8e Pri: 1 Type: Auto
State: FULL DR: 0.0.0.0 BDR: 0.0.0.0 Dead Time: 00:00:31
Options: 0x13 (-|R|-|-|E|V6) Opaque LSA: No
Restart Helper Status: Off
Last Restart Helper Exit Reason: None
BFD Session State: None
```

History

This command was first available in ExtremeXOS 21.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ospfv3 virtual-link timer

```
configure ospfv3 virtual-link {routerid} router_identifier {area}
  area_identifier timer {retransmit-interval} retransmit_interval
  {transit-delay} transit_delay {hello-interval} hello_interval {dead-
interval} dead_interval
```

Description

Configures the timers for a virtual link.

Syntax Description

<i>router_identifier</i>	Specifies the router ID of the other end of the link.
<i>area_identifier</i>	Specifies an <i>OSPFv3</i> area, a four-byte, dotted decimal number.
<i>retransmit_interval</i>	Specifies the length of time that the router waits before retransmitting an LSA that is not acknowledged. The range is 1 to 3600 seconds.
<i>transit_delay</i>	Specifies the length of time it takes to transmit an LSA packet over the interface. The range is 1 to 3600 seconds.
<i>hello_interval</i>	Specifies the interval at which routers send hello packets. The range is 1 to 65535 seconds.
<i>dead_interval</i>	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. The range is 1 to 65535 seconds.

Default

- Retransmit interval—Default: 5 seconds.
- Transit delay—Default: 1 second.
- Hello interval—Default: 10 seconds.
- Dead interval—Default: 40 seconds.

Usage Guidelines

In OSPFv3, all areas must be connected to a backbone area. If the connection to the backbone is lost, it can be repaired by establishing a virtual link.

The smaller the hello interval, the faster topological changes will be detected, but more routing traffic will ensue.

The setting of the retransmit interval should be conservative, or needless retransmissions will result. The value should be larger for serial lines and virtual links.

The transmit delay value should take into account the transmission and propagation delays for the interface.



Note

The wait interval is not separately configurable. It is always equal to the dead interval.

Example

The following command sets the timers on the virtual link to router 6.6.6.6 transiting area 0.0.0.2:

```
configure ospfv3 virtual-link 6.6.6.6 area 0.0.0.2 timer 10 transit-delay 1
hello-interval 20 dead-interval 200
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [*Switch Engine 32.3 Feature License Requirements*](#) document.

NEW! configure pim add tunnel

```
configure pim {ipv4} add {tunnel} [tunnel_name | tunnel all] {sparse |
dense} {passive}
```

Description

Configures an IP interface for PIM tunnel.

Syntax Description

ipv4	Specifies the IPv4 address family.
add	Specifies to add PIM to a VLAN.
tunnel	Specifies the tunnel name.
<i>tunnel_name</i>	Specifies to configure PIM information for interface.
all	Specifies to configure PIM information on all Layer 3 interfaces.
dense	Specifies PIM dense mode (PIM-DM). (Default mode.)

sparse	Specifies PIM sparse mode (PIM-SM).
passive	Specifies a passive interface.

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.

Passive interfaces are host only interfaces that allow a multicast stream from other VLANs to be forwarded to edge hosts. Since they do not peer with other PIM routers, you should not connect a multicast router to a passive interface.

In order for the interface to participate in PIM, PIM must be globally enabled on the switch using the following command: `enable pim`

Example

The following example enables PIM-DM multicast routing on tunnel accounting:

```
configure pim add tunnel accounting dense
```

History

This command was first available in ExtremeXOS 32.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure pim add vlan

```
configure pim {ipv4 | ipv6} add vlan [vlan-name | all] {dense | sparse}
{passive}
```

Description

Configures an IP interface for PIM.

Syntax Description

ipv4	Specifies the IPv4 address family.
ipv6	Specifies the IPv6 address family.
<i>vlan-name</i>	Specifies a <u>VLAN</u> name.
all	Specifies all VLANs.
dense	Specifies PIM dense mode (PIM-DM). (Default mode.)
sparse	Specifies PIM sparse mode (PIM-SM).
passive	Specifies a passive interface.

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.

Passive interfaces are host only interfaces that allow a multicast stream from other VLANs to be forwarded to edge hosts. Since they do not peer with other PIM routers, you should not connect a multicast router to a passive interface.

In order for the interface to participate in PIM, PIM must be globally enabled on the switch using the following command: `enable pim`

Example

The following example enables PIM-DM multicast routing on VLAN accounting:

```
configure pim add vlan accounting dense
```

History

This command was first available in ExtremeXOS 10.1.

The **passive** option was added in ExtremeXOS 11.1.

The **IPv4** and **IPv6** options were added in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure pim anycast-rp

```
configure pim {ipv4 | ipv6} anycast-rp ip_address [policy | none]
```

Description

Configures or removes a rendezvous point (RP) for Anycast RP using PIM (*RFC 4610*).

Syntax Description

ipv4	Specifies IPv4 address family.
ipv6	Specifies IPv6 address family.
anycast-rp	Specifies configuring Anycast RP.
<i>ip_address</i>	Specifies the Anycast RP address.
<i>policy</i>	Specifies the policy file having a list of IP addresses of peer RP nodes. These IP addresses should be specified using NLRI keyword.
none	Specifies removing an Anycast RP and associated policy containing peer information.

Default

N/A.

Usage Guidelines

The Anycast RP using PIM feature provides fast convergence when RP routers fail using PIM protocol without using the source discovery protocol Multicast Source Discovery Protocol (MSDP) for both IPv4 and IPv6 address families.

To view Anycast RP using PIM information, use the `show pim {ipv4 | ipv6} anycast-rp {ip_address}` command.

Example

The following example specifies the router at IP address 10.45.7.12 as the Anycast RP with `policy_file` as the policy file with the list of peer Anycast RP nodes:

```
# configure pim ipv4 anycast-rp 10.45.7.12 policy_file
```

With a `policy_file` of:

```
entry policy1 {  
  if match any{  
    nlri 10.10.10.1/32;  
    nlri 20.20.20.1/32;  
    nlri 30.30.30.1/32;  
  }  
  then {  
    permit;  
  }  
}
```



```
}  
}
```

The following example removes the RP router at IP address 10.45.7.12

```
# configure pim ipv4 anycast-rp 10.45.7.12 none
```

History

This command was first available in ExtremeXOS 31.1.

Platform Availability

This command is available on platforms that support the appropriate license for the PIM feature. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see the [Switch Engine 32.3 Feature License Requirements](#).

configure pim border

```
configure pim {ipv4 | ipv6} [{vlan} vlan_name] border
```

Description

Configures a PIM *VLAN* as a border VLAN, which is used to demarcate a PIM domain when using *MSDP*.

Syntax Description

ipv4	Configures a PIM timer on IPv4 router interfaces.
ipv6	Configures a PIM timer on IPv6 router interfaces.
<i>vlan_name</i>	Specifies a VLAN name.
border	Interface is domain border.

Default

None.

Usage Guidelines

MSDP is used to connect multiple multicast routing domains. A PIM-SM domain is created by limiting the reach of PIM BSR advertisements. When a border VLAN is configured, PIM BSR advertisements are not forwarded out of the PIM VLAN.

Example

The following example configures a PIM border on a VLAN called "vlan_border":

```
configure pim vlan_border border
```

History

This command was first available in ExtremeXOS 10.1.

The **ipv4** and **ipv6** keywords were added giving an option to support this functionality in IPv6 as well in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see the *ExtremeXOS User Guide*.

configure pim cbsr

```
configure pim cbsr {ipv4 | ipv6} [{vlan} vlan_name {priority [0-255]} | none]
```

Description

Configures a candidate bootstrap router for PIM sparse-mode operation.

Syntax Description

ipv4	Specifies the IPv4 address family.
ipv6	Specifies the IPv6 address family.
<i>vlan_name</i>	Specifies a <u>VLAN</u> name.
<i>priority</i>	Specifies a priority setting. The range is 0 - 255.
none	Deletes a CBSR.

Default

The default setting for priority is 0, and indicates the lowest priority.

Usage Guidelines

The VLAN specified for CBSR must have PIM enabled for it to take effect. After PIM is enabled, CBSRs advertise themselves in the PIM domain. A bootstrap router (BSR) is elected among all the candidates based on CBSR priority. To break the tie among routers with the same priority setting, the router with the numerically higher IP address is chosen.

An ExtremeXOS switch can support up to 145 RPs per group when it is configured as a PIM BSR (bootstrap router). If more than 145 RPs are configured for a single group, the BSR ignores the group and does not advertise the RPs. Non-BSR switches can process more than 145 RPs in the BSR message.

Example

The following example configures a candidate bootstrap router on the VLAN accounting:

```
configure pim cbsr vlan accounting 30
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure pim crp static

```
configure pim {ipv4 | ipv6} crp static ip_address [none | policy]  
                {priority [0-255]}
```

Description

Configures a rendezvous point and its associated groups statically, for PIM sparse mode operation.

Syntax Description

ipv4	Specifies an IPv4 address.
ipv6	Specifies an IPv6 address.
<i>ip_address</i>	Specifies a static CRP address.
none	Deletes the static rendezvous point.
<i>policy</i>	Specifies a policy file name.
<i>priority</i>	Specifies a priority setting. The range is 0–255.

Default

The default setting for priority is 192. Priority value 0 indicates the highest priority.

Usage Guidelines

In 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 for the same group (range).

ExtremeXOS switches support up to 64 static RPs (32 IPv4 and 32 IPv6), and up to 180 groups (group/mask entries) in a single RP policy file. If you configure more than 180 group entries in a single RP policy file, the switch will not process entries added after the first 180.

The policy file contains a list of multicast group addresses served by this RP.

This policy file is not used for filtering purposes. As used with this command, the policy file is just a container for a list of addresses. So a typical policy file used for RP configuration looks a little different from a policy used for other purposes.

If routers have different group-to-RP mappings, due to misconfiguration of the static RP (or any other reason), traffic is disrupted.

Example

The following example statically configures an RP and its associated groups defined in policy file rp-list:

```
configure pim crp static 10.0.3.1 rp-list
```

The following is a sample policy file:

```
entry extremel {
  if match any {
  }
  then { nlri 224.0.0.0/4 ;
  nlri 239.255.0.0/24 ;
  nlri 232.0.0.0/8 ;
  nlri 238.1.0.0/16 ;
  nlri 232.232.0.0/20 ;
  }
}
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure pim crp timer

```
configure pim {ipv4 | ipv6} crp timer crp_adv_interval
```

Description

Configures the candidate rendezvous point advertising interval in PIM sparse mode operation.

Syntax Description

ipv4	Specifies an IPv4 address.
ipv6	Specifies an IPv6 address.
<i>crp_adv_interval</i>	Specifies a candidate rendezvous point advertising interval in seconds. The range is 1 to 1,717,986,918.

Default

The default is 60 seconds.

Usage Guidelines

Increasing this time results in increased convergence time for CRP information to the PIM routers.

Example

The following example configures the candidate rendezvous point advertising interval to 120 seconds:

```
configure pim crp timer 120
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure pim crp vlan

```
configure pim {ipv4 | ipv6} crp vlan vlan_name [none | policy]  
           {priority}
```

Description

Configures the dynamic candidate rendezvous point (CRP) for PIM sparse-mode operation.

Syntax Description

ipv4	Specifies the IPv4 address family.
ipv6	Specifies the IPv6 address family.
<i>vlan_name</i>	Specifies a VLAN name.
none	Specifies to delete a CRP.
<i>policy</i>	Specifies a policy file name.
<i>priority</i>	Specifies a priority setting. The range is 0–255.

Default

The default setting for priority is 192. Priority value 0 indicates the highest priority.

Usage Guidelines

ExtremeXOS switches support up to 50 RPs in a switch, and up to 180 groups (group/mask entries) in a single RP policy file. If you configure more than 180 group entries in single RP policy file, then the switch will not process entries added after first 180.

The policy file contains the list of multicast group addresses serviced by this RP. This set of group addresses are advertised as candidate RPs. Each router then elects the common RP for a group address based on a common algorithm. This group to RP mapping should be consistent on all routers.

This policy file is not used for filtering purposes. As used with this command, the policy file is just a container for a list of addresses. So a typical policy file used for RP configuration looks a little different from a policy used for other purposes. The following is a sample policy file that configures the CRP for the address ranges 239.0.0.0/24 and 232.144.27.0/24:

```
entry extremel {
    if match any {
    }
    then {
        nlri 239.0.0.0/24 ;
        nlri 232.144.27.0/24 ;
    }
}
```

The VLAN specified for a CRP must have PIM configured.

To delete a CRP, use the keyword none as the access policy.

Example

The following example configures the candidate rendezvous point for PIM sparse-mode operation on the VLAN HQ_10_0_3 with the policy rp-list and priority set to 30:

```
configure pim crp HQ_10_0_3 rp-list 30
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

NEW! configure pim delete tunnel

```
configure pim {ipv4} delete tunnel [tunnel_name | tunnel all]
```

Description

Disables PIM on a router interface.

Syntax Description

ipv4	Specifies the IPv4 address family.
delete	Specifies to delete PIM on a VLAN.
tunnel	Specifies the tunnel name.
<i>tunnel_name</i>	Specifies to configure PIM information for interface.
all	Specifies to configure PIM information on all Layer 3 interfaces.

Default

N/A.

Usage Guidelines

Use this command to disable PIM for a specific or all tunnels.

Example

The following example disables PIM on tunnel accounting:

```
configure pim delete tunnel accounting
```

History

This command was first available in ExtremeXOS 32.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure pim delete vlan

```
configure {ipv4 | ipv6} pim delete vlan [vlanname | all]
```

Description

Disables PIM on a router interface.

Syntax Description

ipv4	Specifies the IPv4 address family.
ipv6	Specifies the IPv6 address family.
<i>vlanname</i>	Specifies a <u>VLAN</u> name.
all	Specifies all VLANs.

Default

N/A.

Usage Guidelines

Use this command to disable PIM for a specific or all VLANs.

Example

The following example disables PIM on VLAN accounting:

```
configure pim delete vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure pim dense-neighbor-check

```
configure pim dense-neighbor-check [on | off]
```

Description

This command is used to configure a PIM interface that receives multicast data traffic. It could be either from a source directly connected or from a PIM neighbor. In the second case (from a source not directly connected), if the received interface has no PIM neighbor, the traffic is dropped (default behavior). If you turn off this check, the traffic is processed.

Syntax Description

dense-neighbor-check	Check if multicast traffic is received from PIM neighbor in dense mode.
on	Drop multicast traffic if not received from PIM neighbor (default).
off	Forward multicast traffic even if not received from PIM dense neighbor.

Default

The default is on.

Example

The following example turns on dense neighbor check:

```
configure pim dense-neighbor-check on
```

History

This command was first available in ExtremeXOS 15.1.4.

Platform Availability

This command is available on platforms that support the appropriate license. For more information, see the [Switch Engine 32.3 Feature License Requirements](#).

configure pim dr-priority

```
configure pim {ipv4 | ipv6} [ {vlan} vlan_name | vlan all ] dr-priority  
    priority
```

Description

Configures the designated router (DR) priority that is advertised in PIM hello messages.

Syntax Description

ipv4	IPv4 address family (default).
ipv6	IPv6 address family.
vlan all	Apply to all <u>VLANs</u> .
dr-priority	Designated Router Priority for VLAN.
<i>priority</i>	Priority value for VLAN (default 1). The range is 0-4294967295.

Default

The default setting for **dr-priority** is 1.

Usage Guidelines

The **dr-priority** option allows a network administrator to give preference to a particular router in the DR election process by giving it a numerically larger DR priority. The **dr-priority** option is included in every hello message, even if no DR priority is explicitly configured on that interface. This is necessary because priority-based DR election is only enabled when all neighbors on an interface advertise that they are capable of using the **dr-priority** option.

The DR priority is a 32-bit unsigned number, and the numerically larger priority is always preferred. A router's idea of the current DR on an interface can change when a PIM hello message is received, when a neighbor times out, or when a router's own DR priority changes. If the router becomes the

DR or ceases to be the DR, this will normally cause the DR register state machine to change states. Subsequent actions are determined by that state machine. The DR election process on interface is as follows:

- If any one of the neighbor on the interface is not advertised the DR priority (not DR capable) then DR priority will not be considered for all the neighbors in the circuit, and the primary IP address will be considered for all the neighbors.
- The higher DR priority or higher primary address will be elected as DR.

Example

```
configure pim ipv4 vlan accounting dr-priority 10
```

History

This command was first available in ExtremeXOS 15.3.2.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure pim iproute sharing hash

```
configure pim {ipv4 | ipv6} iproute sharing hash [source | group | source-group | source-group-nexthop]
```

Description

This command is used to configure the PIM *ECMP* hash algorithm.

Syntax Description

ipv4	Specifies the IPv4 address family.
ipv6	Specifies the IPv6 address family.
hash	Configure Hash Algorithm for Equal Cost Multipath Routing.
source	Hash for route sharing is based on source address only.
group	Hash for route sharing is based on group address only.
source-group	Hash for route sharing is based on source and group addresses.
source-group-nexthop	Hash for route sharing is based on source, group, and next hop addresses (default).

Default

Source-group-nexthop.

Usage Guidelines

Use this command to modify the hash algorithm used by PIM for path selection.

Example

The following command configures the PIM ECMP hash algorithm based on source-group-nexthop:

```
configure pim ipv6 iproute sharing hash source-group-nexthop
```

History

This command was first available in ExtremeXOS 15.3.2.

Platform Availability

This command is available on platforms that support the appropriate license.

configure pim register-policy

```
configure pim {ipv4 | ipv6} register-policy [rp_policy_name | none]
```

Description

Configures the register filter at the First Hop Router (FHR). This is the router to which the multicast source is connected to.

Syntax Description

ipv4	Specifies the IPv4 address family.
ipv6	Specifies the IPv6 address family.
<i>rp_policy_name</i>	Specifies the Policy File for Register filter.
none	Unconfigures the configured FHR Register filter.

Default

IPv4.

Usage Guidelines

Use this command to add or remove a First Hop Router Register Filter policy.

Example

The following example configures an IPv4 register policy named "entry_policy" at the FHR:

```
configure pim ipv4 register-policy entry_policy
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure pim register-policy rp

```
configure pim {ipv4 | ipv6} register-policy rp [rp_policy_name | none]
```

Description

Configures the register filter at the Rendezvous Point.

Syntax Description

ipv4	Specifies the IPv4 address family.
ipv6	Specifies the IPv6 address family.
<i>rp_policy_name</i>	Specifies the Policy File for RP Register filter.
none	Unconfigures the configured RP Register filter.

Default

N/A.

Usage Guidelines

Use this command to add or remove a Rendezvous Point Register Filter policy.

Example

The following example configures IPv4 register policy named "entry_policy":

```
configure pim ipv4 register-policy rp entry_policy
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure pim register-rate-limit-interval

```
configure pim {ipv4 | ipv6} register-rate-limit-interval interval
```

Description

Configures the initial PIM-SM periodic register rate.

Syntax Description

ipv4	Specifies the IPv4 address family.
ipv6	Specifies the IPv6 address family.
<i>interval</i>	Specifies an interval time in seconds. Range is 0 - 60. Default is 0.

Default

The default interval is 0.

Usage Guidelines

Configuring a non-zero interval time can reduce the CPU load on the first hop switch, in case register stop messages are not received normally.

When a non-zero value is configured, the first hop switch sends a few register messages and then waits for a corresponding register stop from the RP for *time* seconds. The process is repeated until the register stop is received. This command should be used when the (S,G) tree between the first hop router and the RP is not converging quickly.

When the default value is zero in default mode, the switch sends continuous register messages until the register stop is received.

Example

The following example configures the initial PIM register rate limit interval:

```
configure pim register-rate-limit-interval 2
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure pim register-suppress-interval register-probe-interval

```
configure pim {ipv4 | ipv6} register-suppress-interval reg-interval  
                register-probe-interval probe_interval
```

Description

Configures an interval for periodically sending null-registers.

Syntax Description

ipv4	Specifies the IPv4 address family.
ipv6	Specifies the IPv6 address family.
<i>reg-interval</i>	Specifies an interval time in seconds. Range is 30 - 200 seconds. Default is 60.
<i>probe-interval</i>	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). A response to the null register is expected within register probe interval. By specifying a larger interval, a CPU peak load can be avoided because the null-registers are generated less frequently. The register probe time should be less than half of the register suppress time, for best results.

Example

The following example configures the register suppress interval and register probe time:

```
configure pim register-suppress-interval 90 register-probe time 10
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure pim snooping sgrpt-prune

```
configure pim snooping sgrpt-prune [accept | drop]
```

Description

Configures <S,G,RPT> prune messages processing by PIM Snooping.

Syntax Description

accept	<S,G,RPT> prune messages are processed.
drop	<S,G,RPT> prune messages are not processed.

Default

Default configuration is accept.

Usage Guidelines

Use this command when it is desirable to disable PIM <S,G,RPT> prune messages processing by PIM Snooping.

Example

The following example disables <S,G,RPT> prune messages processing by PIM Snooping:

```
configure pim snooping sgrpt-prune drop
```

History

This command was first available in ExtremeXOS 15.6.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see the Feature License Requirements document.

configure pim shutdown-priority

```
configure pim {ipv4 | ipv6} [ {vlan} vlan_name | vlan all ] shutdown-  
priority number
```

Description

Configures the priority for out of memory shutdown.

Syntax Description

ipv4	Configures a PIM timer on IPv4 router interfaces.
ipv6	Configures a PIM timer on IPv6 router interfaces.
<i>vlan_name</i>	Specifies a VLAN name.
all	Specifies all VLANs.
<i>number</i>	Priority for VLAN range is [0 - 65535].

Default

IPv4.

Usage Guidelines

None.

Example

The following example configures the shutdown priority for VLAN 36:

```
config pim vlan v36 shutdown-priority 22
```

History

This command was first available in ExtremeXOS 12.4.

The **ipv4** and **ipv6** keywords were added giving an option to support this functionality in IPv6 as well in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure pim spt-threshold

```
configure pim {ipv4 | ipv6} spt_threshold [infinity | 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. When infinity option is configured on First

Hop Routers or Intermediary Routers, SPT switching is disabled. Traffic forwarding will be performed based on RPT paths only.

Syntax Description

ipv4	Specifies the IPv4 address family.
ipv6	Specifies the IPv6 address family.
infinity	Disables Shortest Path Tree (SPT) switching on Last Hop or Intermediary routers.
<i>leaf-threshold</i>	Specifies the rate of traffic per (s,g,v) group in kbps for the last hop. Range is 0 - 4194303.
<i>rp_threshold</i>	Specifies an RP threshold. Range is 0 - 4194303.

Default

The default setting is 0 for both parameters.

Usage Guidelines

For the best performance, use default value of 0.

Example

The following example changes the threshold for switching to SPT:

```
configure pim spt-threshold 4 16
```

History

This command was first available in ExtremeXOS 10.1.

The **infinity** option was added in ExtremeXOS 15.7.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure pim ssm range

```
configure pim {ipv4 | ipv6} ssm range [default | policy policy-name]
```

Description

Configures the range of multicast addresses for PIM SSM.

Syntax Description

ipv4	Specifies the IPv4 address family.
ipv6	Specifies the IPv6 address family.
default	Specifies the default address range. 232.0.0.0/8 for IPv4 or FF3x::/96 for IPv6.
<i>policy-name</i>	Specifies a policy that defines the SSM address range.

Default

By default, no SSM range is configured. Using this command with the default keyword sets the range to 232.0.0.0/8. To reset the switch to the initial state, use the `unconfigure pim ssm range` command.

Usage Guidelines

Initially, no range is configured for SSM. After a range is configured, you can remove the range with the `unconfigure pim ssm range` command. If you wish to change the PIM SSM range, you must first unconfigure the existing range, and then configure the new range.

SSM requires that hosts use IGMPv3 messages to register to receive multicast group packets. When a range is configured for SSM, any IGMPv2 messages for an address in the range are ignored. Also, any IGMPv3 Exclude messages are ignored.



Note

If a PIM-SSM range is configured, IGMPv2 messages and IGMPv3 exclude messages within the PIM-SSM range are ignored on all IP interfaces, whether or not PIM-SSM is configured on the interfaces.

To specify a range different from the default PIM SSM range, create a policy file. The match statement of the policy file contains the group addresses to be treated as PIM SSM addresses. For example, to specify the PIM SSM address range as 232.0.0.0/8 and 233.0.0.0/8, use the following policy file:

```
Entry extremel {
    if match any {
        nlri 232.0.0.0/8 ;
        nlri 233.0.0.0/8 ;
    }
    then {
        permit ;
    }
}
```

Example

The following example sets the PIM SSM range to 232.0.0.0/8 and 233.0.0.0/8, if the policy file `ssmrange.pol` contains the policy example used above:

```
configure pim ssm range policy ssmrange
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure pim state-refresh timer origination-interval

```
configure pim {ipv4 | ipv6} state-refresh timer origination-interval
interval
```

Description

Configures the interval at which state refresh messages are originated.

Syntax Description

<i>interval</i>	Specifies a refresh interval in seconds. The range is 30–90 seconds.
-----------------	--

Default

60 seconds.

Usage Guidelines

None.

Example

The following example configures the interval to 45 seconds:

```
configure pim state-refresh timer origination-interval 45
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure pim state-refresh timer source-active-timer

```
configure pim {ipv4 | ipv6} state-refresh timer source-active-timer  
    interval
```

Description

Defines how long a multicast source (S,G) is considered active after a packet is received from the source.

Syntax Description

ipv4	Specifies the IPv4 address family.
ipv6	Specifies the IPv6 address family.
<i>interval</i>	Specifies a source-active timer interval in seconds. The range is 90–300 seconds.

Default

210 seconds.

Usage Guidelines

None.

Example

The following example configures the interval to 180 seconds:

```
configure pim state-refresh timer source-active-timer 180
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure pim state-refresh ttl

```
configure pim {ipv4 | ipv6} state-refresh ttl ttlvalue
```

Description

Configures a time-to-live (TTL) value for PIM-DM state refresh messages.

Syntax Description

ipv4	Specifies the IPv4 address family.
ipv6	Specifies the IPv6 address family.
<i>ttl_value</i>	Specifies a TTL value. The range is 1-64.

Default

16.

Usage Guidelines

None.

Example

The following example configures the TTL value for 24:

```
configure pim state-refresh ttl 24
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure pim state-refresh

```
configure pim {ipv4 | ipv6} state-refresh {vlan} [vlan_name | all] [on | off]
```

Description

Enables or disables the PIM-DM state refresh feature on one or all VLANs.

Syntax Description

ipv4	Specifies the IPv4 address family.
ipv6	Specifies the IPv6 address family.

<i>vlan_name</i>	Specifies a VLAN on which to enable or disable the PIM-DM state refresh feature.
on	Enables the PIM-DM state refresh feature on the specified VLANs.
off	Disables the PIM-DM state refresh feature on the specified VLANs.

Default

Disabled.

Usage Guidelines

When this feature is disabled on an interface, the interface behaves as follows:

- State refresh messages are not originated.
- State refresh messages received on the interface are dropped without processing.
- State refresh messages received on other interfaces are not forwarded to the disabled interface.

Example

The following example enables the PIM-DM state refresh feature on VLAN blue:

```
configure pim state-refresh blue on
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

NEW! configure pim timer tunnel

```
configure pim {ipv4} timer hello_interval jp_interval [{tunnel}
    tunnel_name | tunnel all]
```

Description

Configures the global PIM timers on the specified router interfaces.

Syntax Description

ipv4	Specifies the IPv4 address family.
timer	Specifies to configure the timer.

<i>hello_interval</i>	Specifies the amount of time before a hello message is sent out by the PIM router. The range is 1–65,535 seconds.
<i>jp_interval</i>	Specifies the join/prune interval. The range is 1–65,535 seconds.
tunnel	Specifies the tunnel name.
<i>tunnel_name</i>	Specifies the name of the PIM interface.
all	Specifies to apply to all interfaces.

Default

- *hello_interval*—30 seconds
- *jp_interval*—60 seconds

Usage Guidelines

These default timers should only be adjusted when excess PIM control packets are observed on the interface.

Example

The following example configures the PIM timers on the tunnel accounting:

```
configure pim timer 150 300 tunnel accounting
```

History

This command was first available in ExtremeXOS 32.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure pim timer vlan

```
configure pim {ipv4 | ipv6} timer hello_interval jp_interval [{vlan}  
  vlan_name | vlan all]
```

Description

Configures the global PIM timers on the specified router interfaces.

Syntax Description

ipv4	Specifies the IPv4 address family.
ipv6	Specifies the IPv6 address family.

<i>hello_interval</i>	Specifies the amount of time before a hello message is sent out by the PIM router. The range is 1–65,535 seconds.
<i>jp_interval</i>	Specifies the join/prune interval. The range is 1–65,535 seconds.
<i>vlan_name</i>	Specifies a <i>VLAN</i> name.
all	Specifies all VLANs.

Default

- *hello_interval*—30 seconds
- *jp_interval*—60 seconds

Usage Guidelines

These default timers should only be adjusted when excess PIM control packets are observed on the interface.

Example

The following example configures the PIM timers on the VLAN accounting:

```
configure pim timer 150 300 vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

NEW! configure pim tunnel dr-priority

```
configure pim {ipv4} [{tunnel} tunnel_name | tunnel all] dr-priority
priority
```

Description

Configures the tunnel designated router (DR) priority that is advertised in PIM hello messages.

Syntax Description

ipv4	IPv4 address family (default).
tunnel	Specifies the tunnel name.
<i>tunnel_name</i>	Specifies the name of the PIM interface.

dr-priority	Designated Router Priority for VLAN.
<i>priority</i>	Priority value for VLAN (default 1). The range is 0–4294967295.

Default

The default setting for **dr-priority** is 1.

Usage Guidelines

The **dr-priority** option allows a network administrator to give preference to a particular router in the DR election process by giving it a numerically larger DR priority. The **dr-priority** option is included in every hello message, even if no DR priority is explicitly configured on that interface. This is necessary because priority-based DR election is only enabled when all neighbors on an interface advertise that they are capable of using the **dr-priority** option.

The DR priority is a 32-bit unsigned number, and the numerically larger priority is always preferred. A router's idea of the current DR on an interface can change when a PIM hello message is received, when a neighbor times out, or when a router's own DR priority changes. If the router becomes the DR or ceases to be the DR, this will normally cause the DR register state machine to change states. Subsequent actions are determined by that state machine. The DR election process on interface is as follows:

- If any one of the neighbor on the interface is not advertised the DR priority (not DR capable) then DR priority will not be considered for all the neighbors in the circuit, and the primary IP address will be considered for all the neighbors.
- The higher DR priority or higher primary address will be elected as DR.

Example

```
configure pim ipv4 tunnel accounting dr-priority 10
```

History

This command was first available in ExtremeXOS 32.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure pim vlan trusted-gateway

```
configure pim {ipv4 | ipv6} [{vlan} vlan_name] trusted-gateway [policy | none]
```

Description

Configures a trusted neighbor policy.

Syntax Description

ipv4	Specifies the IPv4 address family.
ipv6	Specifies the IPv6 address family.
<i>vlan_name</i>	Specifies a VLAN name.
<i>policy</i>	Specifies a policy file name.
none	Specifies no policy file, so all gateways are trusted.

Default

No policy file, so all gateways are trusted.

Usage Guidelines

Because PIM leverages the unicast routing capability that is already present in the switch, the access policy capabilities are, by nature, different. When the PIM protocol is used for routing IP multicast traffic, the switch can be configured to use a policy file to determine trusted PIM router neighbors for the VLAN on the switch running PIM. This is a security feature for the PIM interface.

Example

The following example configures a trusted neighbor policy on the VLAN backbone using the policy "nointernet":

```
configure pim vlan backbone trusted-gateway nointernet
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure policy access-list

```
configure policy access-list [rule-precedence [list_dot_rule [after  
  member_rule | before member_rule | first | last ] ] ]
```

Description

Adds rules and configures the rule precedence list for an access-list.

Syntax Description

access-list	Configures access-list rule model.
rule-precedence	Specifies modifying a rule's precedence in the access-list.
<i>list_dot_rule</i>	Specifies the access-list name and rule name in the format <i>list_name.rule_name</i> .
after	Moves the rule after an existing entry.
before	Moves the rule before an existing entry.
<i>member_rule</i>	Specifies the access-list name and rule name in format <i>list_name.rule_name</i> .
first	Makes the rule the first.
last	Makes the rule the last.

Default

N/A.

Usage Guidelines

An access-list always contains at least one rule and is not active or programmed until it is assigned to a profile. Assigning a different profile ID to an access-list that already has one overwrites the current value. Setting the profile ID to "none" removes the access-list from the active/programmed rules. A profile ID can only be assigned to an access-list, and not per rule, so the *list_name* must only contain an access-list and not a *list_dot_rule* value.

Example

The following example places the access-list "ACL1.ace3" before "ACL1.ace1":

```
# configure policy access-list rule-precedence ACL1.ace3 before ACL1.ace1
```

History

This command was first available in ExtremeXOS 30.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure policy autoclear

```
configure policy autoclear {interval interval}
```

Description

Sets the interval at which the switch automatically clears rule usage statistics.

Syntax Description

autoclear	Designates setting the parameters for auto-clearing the policy rule usage statistics.
interval	Designates setting the interval when the switch automatically clears rule usage. Default is 0 (statistics are not automatically cleared).
<i>interval</i>	Sets the value for the interval in minutes when the switch automatically clears rule usage. Range is 0 to 65,535.

Default

By default, the **autoclear** interval is 0, which means that statistics are not automatically cleared.

Usage Guidelines

If you have configured Syslog and/or trap actions to notify you when a policy rule is used by using the following command: `configure policy rule profile_index [{app-signature group group name name} | ether ether | icmp6type icmp6type | icmptype icmptype | ip6dest ip6dest | ipdestsocket ipdestsocket | ipfrag | ipproto ipproto | ipsourcesocket ipsourcesocket | iptos iptos | ipttl ipttl | macdest macdest | macsource macsource | port port | tcpdestportIP tcpdestportIP | tcpsourceportIP tcpsourceportIP | udpdestportIP udpdestportIP | udpsourceportIP udpsourceportIP] {mask mask } {port-string [port_string | all]} {storage-type [non-volatile | volatile]} {drop | forward} {syslog syslog} {trap trap} {cos cos } {mirror-destination control_index} {clear-mirror}`, this command allows you to set the interval when these statistics will be cleared.

To view the auto-clear interval, use the following command:

```
show policy autoclear interval
```

Example

The following example sets the interval for automatically clearing rule usage statistics to 1 minute:

```
# configure policy autoclear interval 1
```

History

This command was available in ExtremeXOS 30.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure policy app-signature group name pattern

```
configure policy app-signature group group name [add | delete]  
    pattern_list
```

Description

Configures a user-defined policy application signature.

Syntax Description

app-signature	Configures application signature specific settings.
group	Configures application signature group-specific settings.
<i>group</i>	Specifies the group name.
name	Configures application signature display name-specific settings.
<i>name</i>	Specifies the display name assigned to the application signature. Maximum of 32 characters. To see name choices, use the <code>show policy app-signature group {<i>group</i> {name <i>name</i>}} {built-in custom {detail} detail}</code> command.
add	Adds patterns to the display name.
delete	Removes patterns from the display name.
<i>pattern_list</i>	Specifies a list of strings enclosed in quotes used to identify the application, each separated by a space. Maximum of 255 characters.

Default

N/A.

Usage Guidelines

The application signature groups are built-in and additional ones cannot be created. There are built-in values for application signature names, which cannot be modified or deleted.

Example

The following example for the group name "E-commerce" and application signature name "Warehouse" adds the patterns "bjs.com", "costco.com", and "samsclub.com":

```
# configure policy app-signature group "E-commerce" name Warehouse add "bjs.com  
costco.com samsclub.com"
```

History

This command was first available in ExtremeXOS 30.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure policy app-signature minimum-ttl

```
configure policy app-signature minimum-ttl [none | 1 | 5 | 10]
```

Description

Configures a minimum time-to-live (TTL) value for Layer 7 policy/application signature.

Syntax Description

app-signature	Specifies configuring application signature settings.
minimum-ttl	Specifies setting override to low DNS-reply TTL values with a minimum value. The default is none, which specifies not overriding the TTL values.
none	Specifies not overriding DNS-reply TTL values (default).
1	Specifies a minimum TTL of 1 minute.
5	Specifies a minimum TTL of 5 minutes.
10	Specifies a minimum TTL of 10 minutes.

Default

By default, the DNS-reply TTL values are not overridden (**none**).

Usage Guidelines

To view the TTL minimum value set by this command, use the `show policy app-signature` command.

Example

The following example sets a minimum TTL of 5 minutes:

```
# configure policy app-signature minimum-ttl 5
```

History

This command was first available in ExtremeXOS 30.5.

Limitations

The ExtremeSwitching 5520 series switch does not support Layer 7 policy (DNS).

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure policy captive-portal

```
configure policy captive-portal web-redirect redirect_index server
  server_id {url redirect_url} {status}
```

Description

This command configures a captive portal server's HTTP redirect URL and its status.

Syntax Description

web-redirect	Configures web-redirect.
<i>redirect_index</i>	Configures a web redirect index (range = 1-10).
server	Configures a server for the web redirect index.
<i>server_id</i>	Sets the server ID to use (range = 1-2).
url	Configures captive portal server absolute URL.
<i>redirect_url</i>	Sets HTTP/HTTPS URL that users are redirected to <code>http(s)://<IPv4Address or Hostname>:<L4Port>/<Path></code> Where <code>IPv4Address</code> or <code>Hostname</code> is the IPv4 address or hostname of the captive portal server (DNS server needs to be configured on the device). <code>L4Port</code> by default is 80. Should be provided with the value on which the captive portal web-server is running.
<i>status</i>	Captive portal server status: "enable" or "disable" (default is disable).

Default

By default, captive portal server status is disabled.

Example

The following example configures and enables the URL for a particular captive portal server (index 2) in web-redirect (index 1):

```
configure policy captive-portal web-redirect 1 server 2 url http://192.168.1.1:80/static/
index.jsp enable
```

History

This command was first available in ExtremeXOS 22.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure policy captive-portal listening

```
configure policy captive-portal listening socket_list
```

Description

This command configures which L4 listening ports (sockets) are redirected when a captive portal web-redirect is defined on a policy profile.

Syntax Description

listening	Configures captive portal HTTP listening ports (up to three L4 ports).
<i>socket_list</i>	List of L4 ports on which to listen (1-65,535) separated by commas (for example: 80,8080,2000).

Default

N/A

Usage Guidelines

You can configure a maximum of three L4 listening ports.

Example

The following example configures two L4 listening ports 80 and 8080 to be redirected by captive portal:

```
configure policy captive-portal listening 80,8080
```

The following example add one more L4 listening port 2000:

```
configure policy captive-portal listening 2000
```

The following example tries to apply a fourth listening port 5000. This fails because you can only have three listening ports configured:

```
configure policy captive-portal listening 5000
ERROR: Unable to add 5000. Only 0 remaining socket(s) available.
```

History

This command was first available in ExtremeXOS 22.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure policy captive-portal rule-use

```
configure policy captive-portal rule-use [reserved | unreserved]
```

Description

Configures whether or not captive portal *ACL* rules are programmed within the reserved space for ONEPolicy.

Syntax Description

captive-portal	Configures captive portal elements.
rule-use	Configures captive portal rule use.
reserved	Configures captive portal to program rules in the space reserved by resource-profile configuration at the expense of IPv4 group rules.
unreserved	Configures captive portal to program rules outside the space reserved by resource-profile configuration (default).

Default

By default, captive portal rules are programmed outside of the reserved space for ONEPolicy.

Usage Guidelines

If not specified to do otherwise, ONEPolicy programs its captive portal-related rules outside of the reserved ACL rule space for ONEPolicy (unreserved). This results in additional ACL slice usage. This command enables you to specify that these rules are programmed within the already reserved ACL rule space at the expense of IPv4 rule capacity (reserved).

To view the selection for this command, use the `show policy captive-portal {web-redirect {redirect_index | all} | listening | rule-use}` command with the **rule-use** option.

Example

The following example confines captive portal ACL rules to the reserved space for ONEPolicy:

```
# configure policy captive-portal rule-use reserved
```

History

This command was first available in ExtremeXOS 30.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure policy convergence-endpoint

```
configure policy convergence-endpoint [enable | disable]
```

Description

This command globally enables or disables Convergence End Point (CEP) for ONEPolicy.

Syntax Description

enable	Enables CEP for ONEPolicy.
disable	Disables CEP for ONEPolicy.

Default

By default CEP is disabled.

Usage Guidelines

This feature requires that ONEPolicy is enabled on the switch (see [enable policy](#) on page 2277).

Example

The following example enables CEP on the switch:

```
# configure policy convergence-endpoint enable
```

History

This command was first available in ExtremeXOS 22.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure policy convergence-endpoint clear

```
configure policy convergence-endpoint clear ports [port_list | all]
```

Description

This command clears all existing Convergence End Point (CEP) connections per port.

Syntax Description

ports	Specify ports to configure.
<i>port_list</i>	Designates which ports to clear CEP connections from.

Default

N/A

Example

The following example clears CEP connections from port 3:

```
# configure policy convergence-endpoint clear ports 3
```

History

This command was first available in ExtremeXOS 22.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure policy convergence-endpoint index

```
configure policy convergence-endpoint index index [cisco | lldp-med]
```

Description

This command sets a global default policy index for a Convergence End Point (CEP) detection type. This policy is applied when a phone of the specified type is detected on a port.

Syntax Description

<i>index</i>	The policy index to apply. Use 0 to clear an index. Note: After CEP devices are mapped to a profile, changing the index value to "0" or to some other policy profile name, the existing CEP connections are still be mapped to the old profile that was configured initially when the CEP devices were detected. To force a refresh of existing detected devices, disable, and then enable, CEP (see configure policy convergence-endpoint on page 1118) or disable, and then enable, the port(s) (see disable port on page 1953 and enable port on page 2278).
cisco	Specifies Cisco type CEP.
lldp-med	Specifies <u>LLDP</u> -MED type CEP.

Default

N/A

Usage Guidelines

The corresponding policy must be configured using the policy management commands (for example, [configure policy profile](#) on page 1123).

Example

The following example applies as default the policy associated with index number "12" to Cisco type CEPs.

```
# configure policy convergence-endpoint index 12 cisco
```

History

This command was first available in ExtremeXOS 22.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure policy convergence-endpoint ports

```
configure policy convergence-endpoint ports [<port_list> | all] [cisco | lldp-med] [enable | disable]
```

Description

This command enables or disables a Convergence End Point (CEP) detection type on one or more ports.

Syntax Description

<i>port_list</i>	Specifies ports to configure for CEP detection.
all	Specifies that all ports are configured for CEP detection.
cisco	Selects Cisco type of CEP detection.
lldp-med	Selects <u>LLDP</u> -MED type of CEP detection.
enable	Enables CEP for the provided type.
disable	Disables CEP for the provided type.

Default

By default, CEP detection is disabled on all ports for all types.

Usage Guidelines

This feature requires that ONEPolicy is enabled on the switch (see [enable policy](#) on page 2277).

Example

The following example configures CEP detection for Cisco type on port 3:

```
# configure policy convergence-endpoint ports 3 cisco enable
```

History

This command was first available in ExtremeXOS 22.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure policy invalid action

```
configure policy invalid action [default-policy | drop | forward]
```

Description

This command configures what action is taken for an invalid policy.

Syntax Description

default-policy	Ignore the result and search for the next policy assignment rule.
drop	Block traffic.
forward	Forward traffic as if no policy has been assigned via 802.1D/Q rules.

Default

None.

Example

This example shows how to assign a drop action to invalid policies:

```
X450G2-48t-10G4.4 # configure policy invalid action drop
```

History

This command was first available in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure policy mactable

```
configure policy mactable [response [tunnel | policy | both] | vlan_list
profile_index]
```

Description

Use this command to add entries to the mapping table and to set the map table response state for the switch.

Syntax Description

<i>vlan_list</i>	<u>VLAN</u> ID or range of IDs (1-4,094)
<i>profile_index</i>	Policy ID (1-63).
response	Indicates which attributes to use from RADIUS response.
tunnel	Applies the VLAN-tunnel attribute. VLAN/NSI mappings from RADIUS are used if present. Mappings in policy profile are ignored.
policy	Applies the policy specified in the filter-ID. VLAN/NSI mappings from policy profile are used if present. Mappings in RADIUS response are ignored.
both	An enhanced policy option that applies either all the filter-ID and VLAN tunnel attributes or the policy depending upon whether one or both are present. VLAN/NSI mappings from either RADIUS or policy profile may be used. Mappings in RADIUS response have a higher precedence over policy profile when both contain mappings.

Default

N/A.

Usage Guidelines

The policy response is the default response for the `configure policy mactable` command.

Example

This example adds an entry to the map table that maps VLAN 3 to policy profile 8:

```
configure policy mactable 3 8
```

History

This command was first available in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure policy port

```
configure policy port ports admin-id admin_id
```

Description

This command assigns an administrative rule to a port.

Syntax Description

<i>ports</i>	Port string
admin-id	Policy ID
<i>admin_id</i>	Policy ID (1-63).

Default

N/A.

Usage Guidelines

Use this command to assign an administrative rule to a port.

History

This command was first available in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure policy profile

```
configure policy profile profile_index {name name} {pvid pvid} {pvid-  
status pvid_status} {cos cos} {cos-status cos_status} {egress-  
vlan egress_vlan_list} {forbidden-vlan forbidden_vlan} {untagged-  
vlan untagged_vlan} {append | clear} {tci-overwrite tci_overwrite}  
{precedence [precedence | default]} {auth-override auth_override}  
{nsi [nsi | none]} {web-redirect web_redir_index} {access-list  
[unassigned | list_name | list_name_placeholder]}
```

Description

Creates a policy profile entry.

Syntax Description

<i>profile_index</i>	Policy ID (1-63).
name	Policy profile name.
<i>name</i>	Profile name string 1-64 characters.
pvid-status	PVID status (enable/disable).
pvid	PVID value (0-4,095). Default is 1, which specifies Default VLAN.
cos-status	CoS status (enable/disable).
cos	Class of Service value (0-22).
egress-vlans	Egress <i>VLAN</i> list (1-4094).
forbidden-vlan	Forbidden VLAN list (1-4,094).
untagged-vlans	Untagged VLAN list (1-4,094).
append	Append to one of Egress, Forbidden, Untagged VLAN list.
clear	Clear from one of Egress, Forbidden, Untagged VLAN list.
tci_rewrite	TCI-overwrite status (enable/disable). Note: The ExtremeSwitching 5520 platform does not support TCI-overwrite. Note: With tci-overwrite disabled, you can only add a VLAN to incoming packets that are untagged or priority tagged (priority set, but vlan=0).
auth-override	Configures authentication override using a port profile attribute. No further authentication occurs on the port if enabled.
<i>auth_override</i>	Authentication override status: "enable" or "disable". Default is disabled.
precedence	Specifies setting the policy classification rule precedence. Note: You cannot set a precedence if the rule model is set for ACL Style Policy (access-list). To set the rule model, use the command <code>configure policy rule-model [access-list hierarchical]</code> .
<i>precedence</i>	Sets the rule precedence (for example: 1-2, 10, 12-18, 20-23, 25, 31). To see the supported rules, use <code>show policy profile {all profile_index} {detail}</code> .
default	Sets the default rule precedence, rather than a custom one (1-2, 10, 12-19, 23, 20-22, 25, 31).
web-redirect	Configures web-redirect.
<i>web_redir_index</i>	Configures a web redirect index (range = 1-10). Default is 0, which is disabled.
nsi	Network Service Identifier. For Fabric Attach and VXLAN (VNI = NSI), provides a mechanism to apply the VLAN/NSI mappings in policy using a profile-based attribute.
<i>nsi</i>	NSI 24-bit value ranging from 1 to 16,777,215.

none	No NSI for the VLAN (default).
access-list	Designates assigning an access list to this profile.
unassigned	Removes an assigned access list (default).
<i>list_name</i>	Selects the access list name to assign to this profile. Type the access-list name as shown in the provided list.
<i>list_name_placeholder</i>	Allows you to provide an access-list name that does not currently exist to assign to this profile.

Default

If optional parameters are not specified, none are applied.

Web direct is disabled by default.

The default for NSI is none.

If no PVID value is given, the default is 1 (Default VLAN).

If you do not set a policy classification rule precedence, the default order is used (1–2, 10, 12–19, 23, 20–22, 25, 31).

By default, not access list is assigned to a profile.

Usage Guidelines

Use this command to create a policy profile entry.

Example

This example shows how to create a policy profile 1 named "netadmin" with PVID override enabled for PVID 10, and Class-of-Service override enabled for CoS 5. This profile can use VLAN 10 for untagged egress:

```
# configure policy profile 1 name netadmin pvid-status enable pvid 10 cos-status enable  
cos 5 untagged-vlans 10
```

History

This command was first available in ExtremeXOS 16.1.

The authentication override parameter was added in ExtremeXOS 22.2.

The NSI keyword was added in ExtremeXOS 22.5.

Policy classification rule precedence re-ordering was added in ExtremeXOS 30.2.

Access list capability was added in ExtremeXOS 30.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure policy resource-profile

```
configure policy resource-profile [default | less-acl [more-ipv4 | more-ipv4-no-ipv6 | more-ipv4-no-l2 | more-ipv4-no-mac-no-ipv6] | more-ipv4-no-mac-no-ipv6-no-l2 | more-ipv4-no-ipv6 | more-ipv4-no-mac-no-ipv6 | more-mac-no-ipv6] {profile-modifier [{no-mac no_mac} {no-ipv4 no_ipv4} {no-ipv6 no_ipv6} {no-l2 no_l2}]}
```

Description

Configures a profile that controls the policy rule resources available for MAC/IPv4/IPv6/L2.

Syntax Description

default	Configure a profile with the default settings.
less-acl	Configure a profile that removes some access list resources to be used for rules.
more-ipv4	Configure a profile that adds IPv4 rules.
more-ipv4-no-ipv6	Configure a profile that adds IPv4 rules at the expense of IPv6 rules.
more-ipv4-no-l2	Configure a profile that adds IPv4 rules at the expense of L2 rules. L2 ether rules are accounted for in the first available space from IPv4, IPv6, or MAC group.
more-ipv4-no-mac-no-ipv6	Configure a profile that adds IPv4 rules at the expense of MAC and IPv6 rules
more-ipv4-no-mac-no-ipv6-no-l2	Configure a profile that adds IPv4 rules at the expense of MAC, IPv6, and L2 rules. L2 ether rules are accounted for in the IPv4 group.
more-mac-no-ipv6	Configure a profile that adds MAC rules at the expense of IPv6 rules.
profile-modifier	Specifies modifying the current profile settings.
no-mac	Specifies modifying the current profile, which removes all MAC rules.
<i>no_mac</i>	Specifies removing all MAC rules: "enable" or "disable" (default is disabled).
no-ipv4	Specifies modifying the current profile, which removes all IPv4 rules.
<i>no_ipv4</i>	Specifies removing all IPv4 rules: "enable" or "disable" (default is disabled).
no-ipv6	Specifies modifying the current profile, which removes all IPv6 rules.
<i>no_ipv6</i>	Specifies removing all IPv6 rules: "enable" or "disable" (default is disabled).

no-12	Modify the current profile that removes all L2 rules. L2 ether rules are accounted for in the first available space from IPv4, IPv6, or MAC group.
<i>no_12</i>	Modifier that removes all L2 rules: enable or disable (default is disabled).

Default

By default, the profile modifier is none.

By default, the profile modifier **no-12** is disabled.

Usage Guidelines

You cannot configure the system to use a new resource profiles while policy is enabled. You must disable policy first.

You cannot configure the system to use a new resource-profile where the profile does not fit with existing defined rules. An error message similar to the following appears:

```
Current IPv6 rule usage 1 is higher than max value 0 supplied by profile more-mac-no-ipv6
```

Example

The following example configures the system to use the resource settings of more-ipv4-no-ipv6:

```
configure policy resource-profile more-ipv4-no-ipv6
```

History

This command was first available in ExtremeXOS 22.1.

Profile modification ability was added in ExtremeXOS 22.4.

The profiles **more-ipv4-no-mac-no-ipv6** and **less-acl-more-ipv4-no-mac-no-ipv6** were added in ExtremeXOS 22.4.

The profiles **more-ipv4-no-12** and **more-ipv4-no-mac-no-ipv6-no-12**, and profile modifier **no-12** were added in ExtremeXOS 30.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure policy rule

```
configure policy rule profile_index [{app-signature group group name
  name} | ether ether | icmp6type icmp6type | icmptype icmptype |
  ip6dest ip6dest | ipdestsocket ipdestsocket | ipfrag | ipproto ipproto
  | ipsourcesocket ipsourcesocket | iptos iptos | ipttl ipttl |
```

```

macdest macdest | macsource macsource | port port | tcpdestportIP
tcpdestportIP | tcpsourceportIP tcpsourceportIP | udpdestportIP
udpdestportIP | udpsourceportIP udpsourceportIP ] {mask mask } {port-
string [ port_string | all]} {storage-type [non-volatile | volatile]}
{drop | forward} {syslog syslog} {trap trap} {cos cos } {mirror-
destination control_index} {clear-mirror}

```

Description

Use this command to assign incoming untagged frames to a specific policy profile and to VLAN or CoS classification rules.

Syntax Description

port	Port string.
<i>port</i>	Port string - (data: 1; mask: 16).
app-signature	Associates an application signature to a policy profile.
group	Associates an application signature group to a policy profile
<i>group</i>	Specifies the group name.
name	Associates an application signature name to a policy profile.
<i>name</i>	Specifies the display name assigned to the application signature. Maximum of 32 characters. To see name choices, use the <code>show policy app-signature group {group {name name}} {built-in custom {detail} detail}</code> command.
macsource	MAC source address.
<i>macsource</i>	MAC source address - (data: a-b-c-d-e-f; mask: 1-48).
macdest	MAC destination address.
<i>macdest</i>	MAC destination address - (data: a-b-c-d-e-f; mask: 1-48).
ip6dest	IPv6 address.
<i>ip6dest</i>	IPv6 address (data: aaaa::bbbb; mask 1-128).
ipsourcesocket	Source IP address / Source IpSocket.
<i>ipsourcesocket</i>	Source IP address (data: a.b.c.d[:ab (0-65535)[-cd (0-65535)]]; mask: 1-48, 64).
ipdestsocket	Destination IP address / Destination IpSocket.
<i>ipdestsocket</i>	Destination IP address (data: a.b.c.d[:ab (0-65535) [-cd (0-65535)]]; mask: 1-48,64).
ipfrag	IP fragmentation flag.
tcpdestportIP	TCP port dst with optional post-fix IPv4 address.
<i>tcpdestportIP</i>	TCP port dst with optional post-fix IPv4 address - (data: ab[-cd] [:c.d.e.f]); mask: 1-64).
udpdestportIP	UDP port dst with optional post-fix IPv4 address.

<i>udpdestportIP</i>	UDP port dst with optional post-fix IPv4 address - (data: ab[-cd] [:c.d.e.f]); mask: 1-64.
tcpsourceportIP	TCP port src with optional post-fix IPv4 address.
<i>tcpsourceportIP</i>	TCP port src with optional post-fix IPv4 address - (data: ab[-cd] [:c.d.e.f]); mask: 1-64.
udpsourceportIP	UDP port src with optional post-fix IPv4 address.
<i>udpsourceportIP</i>	UDP port src with optional post-fix IPv4 address - (data: ab[-cd] [:c.d.e.f]); mask: 1-64.
ipttl	IP time to live.
<i>ipttl</i>	ipttl IP time to live (data: 0-255 or 0x0-0xFF; mask:1-8).
iptos	IPv4 type of service / IPv6 traffic class field.
<i>iptos</i>	ipproto Protocol field in IP packet - (data: 0-255 or 0x0-0xFF; mask: 1-8).
ipproto	Protocol field in IP packet.
<i>ipproto</i>	Protocol field in IP packet - (data: 0-255 or 0-0xFF; mask: 1-8).
ether	Type field in Ethernet II packet.
<i>ether</i>	Type field in Ethernet II packet - (data: 0-65535 or 0x0-0xFFFF; mask: 1-16).
icmp6type	Specifies type code in ICMPv6 packet.
<i>icmp6type</i>	ICMPv6 type code [(data: 123.456 (dotted-decimal) or AB-CD (dashed-hexadecimal))] mask: 1-16).
icmp6type	Specifies type code in ICMP packet.
<i>icmp6type</i>	ICMP type code (data: a.b; mask: 1-16).
cos	Class of Service [0-255] or -1 for no CoS or forwarding behavior modification is desired
<i>cos</i>	Class of Service [0-255] or -1 for no CoS or forwarding behavior modification is desired.
mirror-destination	Specifies selecting a mirror destination control index.
<i>mirror-destination</i>	Selects the mirror destination control index. Range is 1 to 4.
clear-mirror	Clears mirroring on this rule.
syslog	Specifies setting a Syslog action when rule is used.
<i>syslog</i>	Enable/disable/prohibit Syslog using event Policy.LogRuleHit on first rule use. By default, a Syslog entry only occurs on the first use of the rule. You can change this using the <code>configure policy syslog [machine-readable <i>machine_readable</i> extended-format <i>extended_format</i> every-time <i>every_time</i>]</code> command.
trap	Specifies setting a trap action when rule is first used.
<i>trap</i>	Enable/disable/prohibit trap on first rule use.

Default

- If mask is not specified, all data bits are considered relevant.
- If port-string is not specified, rule is scoped to all ports.
- By default, a Syslog or trap entry only occurs on the first use of the rule.

Usage Guidelines

Classification rules are automatically enabled when created.



Note

ExtremeSwitching X440-G2 and X620 series switches do not support macsource, macdest, or ip6dest classification rule types. Example:

```
# configure policy rule 1 macsource 00-00-00-00-00-01 port-string 3 drop
ERROR: Set failed!
```



Note

The ExtremeSwitching X870 does not support a port-string with the ip6dest classification rule type.

Example

This example shows how to create (and enable) a classification rule to associate with policy number 1. This rule will drop Ethernet II Type 1526 frames:

```
# configure policy rule 1 ether 1526 drop
```

This example shows how to create (and enable) a classification rule to associate with policy profile number 5. This rule specifies that UDP frames from source port 45 will be forwarded:

```
# configure policy rule 5 udpsourceportip 45 forward forward
```

The following example associates the application signature with group "Storage" and name "mike1" to policy rule "2" to block traffic:

```
# configure policy rule 2 app-signature group "Storage" name "mike1" drop
```

History

This command was first available in ExtremeXOS 16.1.

ICMP and ICMPv6 rule types added in ExtremeXOS 22.5.

Applying mirrors to policies and Syslog/trap actions on rule use was added in ExtremeXOS 30.2.

Application signature capability was added in ExtremeXOS 30.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure policy rule admin-profile

```
configure policy rule admin-profile [ macsource macsource | port port ]
    {mask mask } {port-string [port_string | all] } {storage-type [non-
volatile | volatile]} {admin-pid admin_pid }
```

Description

Use this command to assign incoming untagged frames to a specific policy profile and to [VLAN](#) or Class-of-Service classification rules.

Syntax Description

admin-profile	Policy ID of 0.
macsource	MAC source address.
<i>macsource</i>	MAC source address - (data: a-b-c-d-e-f; mask: 1-48).
port	Port string.
<i>port</i>	Port string - (data: 1; mask: 16).
mask	Number of most significant bits to match data value (rule-meaning)
<i>mask</i>	Number of most significant bits to match data value (rule-meaning). Range: 1 - 144.
port-string	Rule port scope.
<i>port-string</i>	Rule port scope.
all	Scope to all ports.
storage-type	Storage type of this rule.
non-volatile	This entry shall be added to non-volatile storage.
volatile	This entry shall be removed from volatile storage.
admin-pid	Policy ID (1-63).
<i>admin-pid</i>	Policy ID (1-63).

Default

- If mask is not specified, all data bits will be considered relevant.
- If port-string is not specified, rule will be scoped to all ports.

Usage Guidelines

Classification rules are automatically enabled when created.

Example

This example shows how to configure classification rule 2 as an administrative profile and assign it to ingress port 1:1:

```
configure policy rule admin-profile port 1:1 port-string 1:1 admin-pid 2
```

History

This command was first available in ExtremeXOS release 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure policy rule-model

```
configure policy rule-model [access-list | hierarchical]
```

Description

Selects the rule model type for configuring policy rules.

Syntax Description

rule-model	Selects a rule model for configuring and ordering policy-based rules.
access-list	Selects access-list rule model, which allows multiple match criteria per rule along with assignable rule ordering within an access-list.
hierarchical	Selects hierarchical rule model, which allows one match criteria per rule and uses the rule type to assign its precedence (default).

Default

The factory default for rule model is hierarchical.

However, if you are upgrading to ExtremeXOS 30.5 or later, and the switch has an existing policy rules configuration, then the rule model remains hierarchical.

Usage Guidelines

To configure rule models, policy must be disabled.

If you change rule models, the configuration of the other rule model is deleted.

Example

The following example sets the rule model to hierarchical:

```
# configure policy rule-model hierarchical
```


History

This command was first available in ExtremeXOS 30.5.

The default rule model was changed from "access list" to "hierarchical" in ExtremeXOS 30.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure policy slices shared

```
configure policy slices shared [{ shared } { l7GuaranteedPercentage
l7GuaranteedPercentage } { dynAclGuaranteedPercentage
dynAclGuaranteedPercentage}]
```

Description

Configures the number of slices used by shared features, such as Layer 7 policy and dynamic [ACL](#).

Syntax Description

slices	Configures look-up stage TCAM resources.
shared	Designates setting the shared lookup stage TCAM resources.
<i>shared</i>	Sets the shared slice value (range is 0-4).
l7GuaranteedPercentage	Designates setting the percentage of shared slice that Layer 7 is guaranteed. Note: The ExtremeSwitching X435 platform does not support Layer 7 policy (DNS).
<i>l7GuaranteedPercentage</i>	Specifies the guaranteed Layer 7 percentage value (range is 0-100).
dynAclGuaranteedPercentage	Designates setting the percentage of shared slice that is dynamic ACL guaranteed.
<i>dynAclGuaranteedPercentage</i>	Specifies the guaranteed dynamic ACL percentage value (range is 0-100).

Default

N/A.

Usage Guidelines

To make changes using this command, you must first disable policy (`disable policy`).

To view selections made by this command, use the `show policy slices` command.

Example

The following example configures policy to use 2 slices for shared features and allocate a guaranteed 40% to Layer 7 and 40% to dynamic ACLs:

```
# configure policy slices shared 2 17GuaranteedPercentage 40 dynAclGuaranteedPercentage 40
```

History

This command was first available in ExtremeXOS 30.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure policy slices tci-overwrite

```
configure policy slices {shared shared} {tci-overwrite slices}
```

Description

Configures the number of slices used by a profile in the look-up stage TCAM resources.

Syntax Description

slices	Configures look-up stage TCAM resources.
tci-overwrite	Configures look-up Stage TCAM resources used by profile with tci-overwrite enabled. Note: The ExtremeSwitching X435 platform does not support TCI-overwrite.
<i>slices</i>	Specifies the number of slices between 0 and 4. The default is 4.
shared	Configures look-up stage TCAM resources.
<i>shared</i>	Specifies the shared slice value (0-4).

Default

By default, the number of slices is 4.

Usage Guidelines

This command only runs if policy is disabled.

This command enables you to allocate only the slice resources necessary and allow the rest to be used outside of policy. In a stack with slots having differing VCAP slice depths, each slot has the number of rules available as follows: `numSlices * (VCAP slice depth)`.

Example

The following example configures policy to use 3 slices with tci-overwrite enabled:

```
# configure policy slices tci-overwrite 3
```

History

This command was first available in ExtremeXOS 30.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure policy syslog

```
configure policy syslog [machine-readable machine_readable | extended-format extended_format | every-time every_time]
```

Description

Sets Syslog parameters for policy rules.

Syntax Description

syslog	Sets Syslog parameters for policy rules.
machine-readable	Sets whether hexadecimal or decimal format is used for Syslog messages.
<i>machine_readable</i>	Sets whether hexadecimal or decimal format is used for Syslog messages: "enable" (= hexadecimal) or "disable" (= decimal). Default is disabled (decimal).
extended-format	Sets whether extended format is used for Syslog messages.
<i>extended_format</i>	Sets whether extended format is used for Syslog messages: "enable" (= extended) or "disable" (= not extended). Default is disabled (not extended).
every-time	Sets whether Syslog messages are sent every time a rule is used (not just first time).
<i>every_time</i>	Sets whether Syslog messages are sent every time a rule is used (not just first time): "disable" or "enable". Default is disabled.

Default

By default, Syslog messages are only sent on first use of a rule.

By default, **extended-format** and **machine-readable** are disabled (not extended and in decimal format).

Usage Guidelines

This command allows you to set parameters for Syslog messages that are sent when a policy rule is used when set up in the command `configure policy rule profile_index [{app-signature group group name name} | ether ether | icmp6type icmp6type | icmptype icmptype | ip6dest ip6dest | ipdestsocket ipdestsocket | ipfrag | ipproto ipproto | ipsourcesocket ipsourcesocket | iptos iptos | ipttl ipttl | macdest macdest | macsource macsource | port port | tcpdestportIP tcpdestportIP | tcpsourceportIP tcpsourceportIP | udpdestportIP udpdestportIP | udpsourceportIP udpsourceportIP] {mask mask } {port-string [port_string | all]} {storage-type [non-volatile | volatile]} {drop | forward} {syslog syslog} {trap trap} {cos cos } {mirror-destination control_index} {clear-mirror} .`

When Syslog messages are configured to be sent every time a rule is used, messages are sent at a maximum rate of once every five seconds.

To view the parameters configured by this command, use the command `show policy syslog {machine-readable} {extended-format} {every-time}`.

Example

The following example sets Syslog messages to be sent every time a rule is used:

```
#configure policy syslog every-time enable
```

History

This command was first available in ExtremeXOS 30.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure policy vlanauthorization

```
configure policy vlanauthorization [enable | disable]
```

Description

This command enables or disables the configuration of VLAN Authorization-specific settings.

Syntax Description

enable	Enable VLAN Authorization.
disable	Disabel VLAN Authorization

Default

N/A.

Usage Guidelines

None.

Example

This example shows how to enable VLAN Authorization:

```
X450G2-48t-10G4.4 # configure policy vlanauthorization enable
```

History

This command was first available in ExtremeXOS 16.1

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure policy vlanauthorization port

```
configure policy vlanauthorization port [ port_list | all ] [{enable |  
disable} {tagged | untagged } ]
```

Description

This command configures VLAN Authorization for a port, port list, or all ports.

Syntax Description

<i>port_list</i>	List of ports to configure.
all	Configure all ports.
enable	Enable VLAN Authorization on port.
disable	Disable VLAN Authorization on port.
tagged	Add port to egress of the VLAN-ID returned.
untagged	Add port to the untagged egress of the VLAN-ID returned.

Default

N/A.

Usage Guidelines

None.

Example

This example shows how to enable VLAN Authorization for port 1:1 for tagged packets:

```
x450G2-48t-10G4.5 # configure policy vlanauthorization port 1:1 enable tagged
```

History

This command was first available in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure port description-string

```
configure ports port_list description-string string
```

Description

Configures a description string setting up to 255 characters.

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
<i>string</i>	Specifies a port description of up to 255 characters per port. You cannot use the following characters: “ ‘ , “<”, “>”, “:”, “<space>”, “&”

Default

None.

Usage Guidelines

Use this command to configure a port description of up to 255 characters per port.

In case that user configures a string longer than 64 chars, the following warning will be displayed:

```
Port description strings longer than 64 chars are only accessible
through SNMP if the following command is issued: configure snmp ifmib
ifalias size extended
```

Some characters are not permitted as they have special meanings. These are: “ ‘ , “<”, “>”, “:”, “<space>”, “&”. The first character should be alphanumeric. This new field is CLI accessible only via “show port info detail” but is also accessible via the [SNMP](#) ifAlias object of IfXTable from IF-MIB (RFC 2233) and the XML API. In order to access the value via SNMP the following command should be issued: configure snmp ifmib ifalias size extended.

Example

The following command configures the port:

```
configure ports 1:3 description-string CorporatePort_123
```

History

This command was available in ExtremeXOS 15.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure port ethertype

```
configure port port_list ethertype {primary | secondary}
```

Description

Assigns the primary or secondary ethertype value to the specified ports.

Syntax Description

<i>port_list</i>	Specifies the list of ports to be configured.
primary	Assigns the primary ethertype value to the specified ports.
secondary	Assigns the secondary ethertype value to the specified ports.

Default

N/A.

Usage Guidelines

None.

Example

The following example configures port 2:1 to use the secondary ethertype:

```
configure port 2:1 ethertype secondary
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure port reflective-relay

```
configure port port reflective-relay [on | off]
```

Description

Enables the direct attach feature on the specified port.

Syntax Description

<i>port</i>	Specifies a single port on which to enable the direct attach feature.
-------------	---

Default

Off.

Usage Guidelines

You should only enable the direct attach feature on ports that directly connect to a VM server running VEPA software.

This feature requires installation of the Direct Attach feature pack. For more information, see the [Switch Engine 32.3 Feature License Requirements](#) document..

Example

The following command enables the direct attach feature on port 2:1:

```
# configure port 2:1 reflective-relay on
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure port shared-packet-buffer

```
configure port port_list shared-packet-buffer [percent | default]
```

Description

Configures the maximum amount of the shared packet buffer to be used by the specified ports.

Syntax Description

<i>port_list</i>	Specifies a list of ports or slots and ports.
<i>percent</i>	Specifies the maximum portion of the shared packet buffer to allot. The range is 0 to 100 percent.



Note

On some platforms, the hardware provides a limited number of settings. In these cases, ranges of percentage values achieve the same setting.



Note

You can view the configured percentage value using the `show ports port-list info detail` command.



Note

You can view the effect of this command using the `show ports port-list buffer` command.

Default

None.

Usage Guidelines

It is possible to overcommit the shared packet buffer using this command.

Example

The following command sets the shared packet buffer for port 1:1 to 50%:

```
configure port 1:1 shared-packet-buffer 50
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ports

```
configure ports {group} port_group [[ add | delete ] port_list ]
```

Description

Creates or deletes a generic port-group name that can be associated with a list of ports.

Syntax Description

group	Named list of ports.
<i>port_group</i>	Port group name.
add	Add ports to port group.
delete	Delete ports from port group.
<i>port_list</i>	Specifies a port list.

Default

N/A.

Usage Guidelines

Use this command to add or delete a generic port-group name to a list of ports.



Note

Because port-groups may be configured for multiple applications, no check is done other than that the values entered are ports. Individual applications handle illegal actions on ports as necessary. QoS commands that use port groups are updated automatically if the ports group is removed or if ports are added or removed from the group.

Example

```
configure ports group testGroup add 1-5
configure ports testGroup delete 3
```

History

This command was first available in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ports auto off

```
configure ports port_list {medium [copper | fiber]} auto off speed speed
duplex [half | full]
```

Description

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

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
medium	Specifies the medium as either copper or fiber. Note: This parameter applies to combo ports..
<i>speed</i>	Specifies the port speed as either 10, 100, 1,000 (1 Gigabit), 2,500 (2.5 Gigabit), 5,000 (5 Gigabit), 10,000 (10 Gigabit), 25,000 (25 Gigabit), 40,000 (40 Gigabit), 50,000 (50 Gigabit), or 100,000 (100 Gigabit) Mbps ports.
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 for 1G and 10G copper ports.

Auto off for 25G, 40G, 50G, and 10,00G ports.

Usage Guidelines

You can manually configure the duplex setting and the speed on 10/100 and 10/100/1000 Mbps and fiber SFP gigabit Ethernet ports.

In general, SFP gigabit Ethernet ports are statically set to 1 Gbps, and their speed cannot be modified. However, there are SFPs supported by Extreme Networks that can have a configured speed:

- 100 FX SFPs, which must have their speed configured to 100 Mbps.
- 100FX/1000LX SFPs, which can be configured at either speed.
- SFP+ optics, must have their speed configured to 10G auto off.

In certain interoperability situations, it is necessary to turn autonegotiation off on a fiber 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 more detailed information about flow control on Extreme Networks devices, see the [Switch Engine 32.3 User Guide](#).

When configuring combination ports you can specify the medium as copper or fiber. If the medium is not specified for combination ports then the configuration is applied to the current primary

medium. The current primary medium is displayed in the Media Primary column of the `show ports configuration` command output.



Note

The keyword **medium** is used to select the configuration medium for combination ports. If the `port_list` contains any non-combination ports, the command is rejected.

When upgrading a switch running ExtremeXOS 12.3 or earlier software to ExtremeXOS 12.4 or later, saved configurations from combo ports (copper or fiber) are applied only to combo ports fiber medium. When downgrading from ExtremeXOS 12.4 or later to ExtremeXOS 12.3 or earlier, saved configurations from combo ports (copper or fiber) are silently ignored. Therefore, you need to reconfigure combo ports during such an upgrade or downgrade.

Example

The following example turns autonegotiation off for port 2 with copper medium and a port speed of 100 Mbps at full duplex:

```
configure ports 2 medium copper auto off speed 100 duplex full
```

History

This command was first available in ExtremeXOS 10.1.

The **medium** parameter was added in ExtremeXOS 12.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ports auto on

```
configure ports port_list {medium [copper|fiber]} auto on [{speed
  speed} {duplex [half | full]}] | [{duplex [half | full]} {speed
  speed}] }
```

Description

Enables autonegotiation for the particular port type.

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
medium	Specifies the medium as either copper or fiber. Note: This parameter applies to combo ports.
<i>speed</i>	Specifies the port speed as either 10, 100, 1,000 (1 Gigabit), 2,500 (2.5 Gigabit), 5,000 (5 Gigabit), 10,000 (10 Gigabit), 25,000 (25 Gigabit), 40,000 (40 Gigabit), 50,000 (50 Gigabit), or 100,000 (100 Gigabit) Mbps ports.

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 for 1G and 10G copper ports.

Auto off for 25G, 40G, 50G, and 10,00G ports.

Usage Guidelines

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

Flow control on gigabit Ethernet ports 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. See the [Switch Engine 32.3 User Guide](#) for more detailed information on flow control on Extreme Networks devices.

When configuring combo ports you can specify the medium as copper or fiber. If the medium is not specified for combination ports then the configuration is applied to the current primary medium. The current primary medium is displayed in the Media Primary column of the [show ports configuration](#) command output.



Note

The keyword medium is used to select the configuration medium for combination ports. If the port_list contains any non-combination ports, the command is rejected.

When upgrading a switch running ExtremeXOS 12.3 or earlier software to ExtremeXOS 12.4 or later, saved configurations from combo ports (copper or fiber) are applied only to combo ports fiber medium. When downgrading from ExtremeXOS 12.4 or later to ExtremeXOS 12.3 or earlier, saved configurations from combo ports (copper or fiber) are silently ignored. Therefore, you need to reconfigure combo ports during such an upgrade or downgrade.

Version 32.2 supports a 2.5G connection between ExtremeSwitching X435 uplink ports and 5520-24X front panel ports. This feature also supports connecting a 5520-24X to another 5520-24X at 2.5G. Use of this feature requires 10G-SR-SFP300M-ET and 10G-LR-SFP10KM-ET 10G transceivers.



Note

For switches that do not support half-duplex (the ExtremeSwitching 5520-12MW-36W multi-rate ports), the copper switch ports must have auto negotiation disabled and full duplex enabled when connecting 10/100/1000 Mbps devices that do not auto negotiate. If the switch attempts and fails to auto negotiate with its partner, it will fail to link up. A non-negotiating connected device must also be manually configured for full duplex or packet loss and port errors will occur each time it detects a collision.



Note

1G auto negotiation is not supported in 5420-YE and 5420-XE ports.

**Note**

10G auto negotiation is supported in 5420 switches while using SFP+ passive copper cables.

**Note**

2.5G support for X435 and 5520-24X is only available with auto-negotiation disabled.

Example

The following example configures the switch to auto-negotiate for port 2, with copper medium at a port speed of 100 Mbps at full duplex:

```
# configure ports 2 medium copper auto on speed 100 duplex full
```

History

This command was first available in ExtremeXOS 10.1.

The **speed** and **duplex** parameters were added in ExtremeXOS 11.6.

The **medium** parameter was added in ExtremeXOS 12.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ports auto-polarity

```
configure ports port_list auto-polarity [off | on]
```

Description

Configures the autopolarity detection feature on the specified Ethernet ports.

Syntax Description

<i>port_list</i>	Specifies one or more ports on the switch.
off	Disables the autopolarity detection feature on the specified ports.
on	Enables the autopolarity detection feature on the specified ports.

Default

Enabled.

Usage Guidelines

This feature applies to only the 10/100/1000 BASE-T ports, and copper medium on combination ports.

When autopolarity is disabled on one or more Ethernet ports, you can verify that status by using the following command:

```
# show ports information detail
```

Example

The following command disables the autopolarity detection feature on ports 5 to 7 on a switch:

```
# configure ports 5-7 auto-polarity off
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ports ddmi

```
configure ports [port_list | all] ddmi [on | off]
```

Description

Enables or disables Digital Diagnostic Monitoring Interface (DDMI).

Syntax Description

<i>port_list</i>	Designates the ports to enable or disable DDMI on.
all	Designates enabling or disabling DDMI on all ports.
ddmi	Designates enabling or disabling DDMI on specified ports.
on	Specifies enabling DDMI on the selected ports (default).
off	Specifies disabling DDMI on the selected ports.

Default

By default, DDMI is enabled.

Usage Guidelines

DDMI provides critical system information about the installed optical modules.

Example

The following example disables DDMI on port 1:

```
# configure ports 1 ddmi off
```

History

This command was first available in ExtremeXOS 30.7.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ports display-string

```
configure ports port_list display-string string
```

Description

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

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
<i>string</i>	Specifies a user-defined display string.

Default

The null string is the default.

Usage Guidelines

The display string can be up to 15 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 information` command.



Note

Do not use a port number as a display string. For example, do not assign the display string “2” to port2.

Example

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

```
configure ports 1 display-string corporate
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ports dot1p

```
configure ports [port_list | all] dot1p dot1p_priority
```

Description

This command configures the default dot1p priority to be used for the internal priority for untagged traffic on the specified port.

Syntax Description

<i>port_list</i>	Specifies a port list.
all	Specifies all ports.
<i>dot1p_priority</i>	Priority number from 0 to 7 to be used for untagged packets.

Default

0.

Usage Guidelines

Use this command to configure the default dot1p priority to be used for the internal priority for untagged traffic on the specified port. This priority is used for untagged frames when dot1p examination is enabled on a port.

History

This command was first available in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ports dwdm channel none

```
configure port all | port_list dwdm channel none
```

Description

Configures the default DWDM channel number.

Syntax Description

all	Specifies all ports.
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

Channel number - 21.

Usage Guidelines

Use this command to configure the default DWDM channel number to the DWDM optical module inserted in the given port. This default channel number of 21 and will be mapped to the appropriate corresponding channel number of the vendor specific channel. If a non-tunable DWDM optic is present, then the DWDM configuration is silently removed from the software.

Example

The following command configures the default DWDM channel 21 on supported port 1:

```
configure port 1 dwdm channel none
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ports dwdm channel

```
configure port all | port_list dwdm channel channel_number
```

Description

Selects the DWDM channel frequency for the selected ports.

Syntax Description

all	Specifies all ports.
<i>port_list</i>	Specifies one or more ports or slots and ports.
<i>channel_number</i>	Specifies the channel number, which corresponds to one of 102 available channel frequencies.

Default

Channel number – 21.

Usage Guidelines

The following table lists the available frequencies and the channel number you must specify to select each frequency.

Table 16: TX Wavelengths and Channel Assignments for the Tunable DWDM XFP/SPF+

TX Wavelength	Channel	TX Wavelength	Channel	TX Wavelength	Channel	TX Wavelength	Channel
1568.77 nm	11	1558.17 nm	24	1547.72 nm	37	1537.40 nm	50
1568.36 nm	1150	1557.77 nm	2450	1547.32 nm	3750	1537.00 nm	5050
1567.95 nm	12	1557.36 nm	25	1546.92 nm	38	1536.61 nm	51
1567.54 nm	1250	1556.96 nm	2550	1546.52 nm	3850	1536.22 nm	5150
1567.13 nm	13	1556.55 nm	26	1546.12 nm	39	1535.82 nm	52
1566.72 nm	1350	1556.15 nm	2650	1545.72 nm	3950	1535.43 nm	5250
1566.31 nm	14	1555.75 nm	27	1545.32 nm	40	1535.04 nm	53
1565.90 nm	1450	1555.34 nm	2750	1544.92 nm	4050	1534.64 nm	5350
1565.50 nm	15	1554.94 nm	28	1544.53 nm	41	1534.25 nm	54
1565.09 nm	1550	1554.54 nm	2850	1544.13 nm	4150	1533.86 nm	5450
1564.68 nm	16	1554.13 nm	29	1543.73 nm	42	1533.47 nm	55
1564.27 nm	1650	1553.73 nm	2950	1543.33 nm	4250	1533.07 nm	5550
1563.86 nm	17	1553.33 nm	30	1542.94 nm	43	1532.68 nm	56
1563.45 nm	1750	1552.93 nm	3050	1542.54 nm	4350	1532.29 nm	5650
1563.05 nm	18	1552.52 nm	31	1542.14 nm	44	1531.90 nm	57
1562.64 nm	1850	1552.12 nm	3150	1541.75 nm	4450	1531.51 nm	5750
1562.23 nm	19	1551.72 nm	32	1541.35 nm	45	1531.12 nm	58
1561.83 nm	1950	1551.32 nm	3250	1540.95 nm	4550	1530.72 nm	5850
1561.42 nm	20	1550.92 nm	33	1540.56 nm	46	1530.33 nm	59
1561.01 nm	2050	1550.52 nm	3350	1540.16 nm	4650	1529.94 nm	5950
1560.61 nm	21	1550.12 nm	34	1539.77 nm	47	1529.55 nm	60
1560.20 nm	2150	1549.72 nm	3450	1539.37 nm	4750	1529.16 nm	6050
1559.79 nm	22	1549.32 nm	35	1538.98 nm	48	1528.77 nm	61
1559.39 nm	2250	1548.91 nm	3550	1538.58 nm	4850	1528.38 nm	6150
1558.98 nm	23	1548.51 nm	36	1538.19 nm	49		
1558.58 nm	2350	1548.11 nm	3650	1537.79 nm	4950		

The supported channel numbers are not contiguous. If you specify a channel number that is not listed in the preceding table, the following error message appears:

```
Error: DWDM Channel configuration failed. Channel number 100 is out of
configurable range. The channel range for the Optical module in port
<port number> is 11 .. 6150.
```

If the optical module in one of the ports in the specified list does not support DWDM, the following error message is displayed:

```
Error: No TDWDM Optics on port <port number>.
```

If the optical module in one of the ports in the specified port list is not an Extreme supported optical module, the following error message is displayed:

```
Error: DWDM Channel configuration failed. Optical module is not Extreme
Networks certified. For DWDM channel configuration, Extreme Network
Certified DWDM module is required.
```

To display the configuration, use the [show ports configuration](#) or the [show ports information](#) detail command.

Example

The following command configures DWDM channel 21 on a modular port 1:1:

```
configure port 1:1 dwdm channel 21
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ports eee

```
configure ports port_list eee [on | off]
```

Description

Enables or disables EEE on the physical layer.

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
on	Specifies that the port advertises to its link partner that it is EEE capable at certain speeds
off	Specifies that the port advertises to its link partner that it is not EEE capable at certain speeds

Default

Off.

Usage Guidelines

Use this command to enable EEE on the switch. The keyword **on** specifies that the port advertises to its link partner that it is EEE capable at certain speeds. If both sides, during auto-negotiation, determine that they both have EEE on and are compatible speed wise, they will determine other parameters (how long it takes to come out of sleep time, how long it takes to wake up) and the link comes up. During periods of non-activity, the link will shut down parts of the port to save energy. This is called LPI for low power idle. When one side sees it must send something, it wakes up the remote and then transmits.

Example

The following example turns the EEE feature on for port 2:

```
config port 2 eee on
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

EEE is supported on the following Extreme Networks platforms:

- ExtremeSwitching 5320, 5420, 5520, and 5720—copper 10/100/1000 and multi-rate ports

configure ports forward-error-correction

```
configure ports port_list forward-error-correction [off | on [c174 | c191]]
```

Description

Enables/disables IEEE Forward Error Correction (FEC) Clause 74 or 91 modes.

Syntax Description

<i>port_list</i>	List of ports to enable/disable FEC modes on.
forward-error-correction	Configures port FEC mode.
off	Disables all FEC modes (default).
on	Enables FEC modes.
c174	Enables/disables FEC IEEE Clause 74.
c191	Enables/disables FEC IEEE Clause 91.

Default

FEC is not enabled by default.

Usage Guidelines

This command allows you to enable/disable Clause 91 or Clause 74 (exclusively) on a per-port basis regardless of speed/type.

FEC gives the receiver the ability to correct errors without requiring a reverse channel to request retransmission of data, but at the cost of a fixed, higher forward channel bandwidth. Some devices require this to interoperate.

Example

The following example enables FEC Clause 91 on port 1:

```
# configure ports 1 forward-error-correction on c191
```

The following example turns off FEC on port 1:

```
# configure ports 1 forward-error-correction off
```

History

This command was first available in ExtremeXOS 22.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ports ingress-filtering

```
configure ports port_list mac-based-vlans ingress-filtering [on|off]
```

Description

Allows users to toggle the ingress filtering setting for MAC-based VLANs.

Syntax Description

<i>port_list</i>	Specifies one or more ports, or slots and ports.
mac-based-vlans	Specifies MAC-based VLANs.
ingress-filtering	Specifies the configuration of ingress filtering when MAC-based VLANs are enabled.
on	Turns on ingress filtering. Enabling ingress filtering prevents VLANs other than those with egress membership from being forwarded.
off	Turns off ingress filtering. Disabling ingress filtering allows the forwarding of packets sourced from any VLAN. This is the default setting.

Default

The default is off.

Usage Guidelines

If the command is executed on a port that does not have MAC-based VLANs enabled, the command will be accepted, but it will not be applied until MAC-Based VLANs are enabled.

Example

The following command enables ingress filtering for ports 1:2:

```
# configure ports 1:2 mac-based-vlans ingress-filtering on
```

History

This command was first available in ExtremeXOS 31.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ports isolation

```
configure ports port_list isolation[on|off]
```

Description

Enables isolation mode on a per-port basis.

Syntax Description

<i>port_list</i>	Specifies one or more ports, or slots and ports.
isolation	Specifies that Isolated ports are not allowed to inter-communicate.
on	Turns on isolation. Isolated ports are not allowed to inter-communicate.
off	Turns off isolation. This is the default setting.

Default

Isolation is off by default.

Usage Guidelines

Use this command to enable isolation mode on a per-port basis. You can issue the command on a single port or on a master port of a load share group. If you issue the command on a non-master port of a load share group the command will fail. When a port load share group is formed, all of the member ports assume the same isolation setting as the master port.

Example

The following command enables isolation mode on ports 2 and 4 on a switch:

```
configure ports 1, 4 isolation on
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ports l2pt profile

```
configure [vlan | vman] vlan_name ports port_list l2pt profile [none | profile_name]
```

Description

Configures L2PT profiles on service interfaces.

Syntax Description

vlan	Specifies the <u>VLAN</u> configuration.
vman	Specifies the VMAN configuration.

<i>vlan_name</i>	Specifies the VLAN name.
ports <i>port_list</i>	Specifies the port and port list separated by a comma (,) or dash (-).
profile	Specifies the L2PT profile for the ports.
none	Specifies that no L2PT profile should be bound to the ports (default).
<i>profile_name</i>	Specifies the L2PT profile to be bound to the ports.

Default

Disabled.

Usage Guidelines

Use this command to configure L2PT profiles on service interfaces.

Example

The following example binds *my_l2pt_prof* with ports 2 and 5 of VMAN cust1:

```
configure vman cust1 ports 2,5 l2pt profile my_l2pt_prof
```

The following example binds *my_l2pt_prof* with ports 2 and 5 of VMAN cust1. Port 5 is not a part of VMAN cust1:

```
configure vman cust1 ports 2,5 l2pt profile my_l2pt_prof
Error: Port 5 is not part of the service.
```

History

This command was first available in ExtremeXOS 15.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ports link-flap-detection action

```
configure ports [port_list | all] link-flap-detection action [add | delete] [{disable-port} {log} {trap}] | all-actions
```

Description

Add or deletes actions (disabling ports, logging events, generating *SNMP* traps) to be taken when excessive link flapping is detected.

Syntax Description

ports	Physical ports.
<i>port_list</i>	List of ports to set link-flap detection actions upon.
all	Sets the configured action for link-flap detection upon all ports in the system.
action	Sets actions to be taken when excessive link flapping is detected.
add	Adds action(s).
delete	Deletes action(s).
disable-port	Disables selected ports if link-flap threshold is exceeded. After a port is disabled, the port either stays down for the configured disable time value (set in the <code>configure ports link-flap-detection interval threshold disable-time</code> command) or can be re-enabled manually using the <code>clear ports link-flap-detection status</code> command.
log	Generates a log event if link-flap threshold is exceeded.
trap	Generates an SNMP trap if link-flap threshold is exceeded.
all-actions	Adds or deletes all the actions.

Default

By default, all actions are turned off.

Example

The following example adds all link-flap actions (disabling ports, logging events, generating SNMP traps) on ports 3-10:

```
configure ports 3-10 link-flap-detection action add all-actions
```

History

This command was first available in ExtremeXOS 22.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ports link-flap-detection interval threshold disable-time

```
configure ports [port_list | all] link-flap-detection [{interval
    [interval | indefinitely]} {threshold threshold} {disable-time
    [disable_time | until-cleared]}]
```

Description

Sets interval, threshold (maximum number of link down events), and disable time values for link-flap detection.

Syntax Description

ports	Physical ports.
<i>port_list</i>	List of ports to activate link-flap detection upon.
all	Sets link-flap detection characteristics on all ports in the system.
interval	Sets time interval for collecting link-flap events.
<i>interval</i>	Interval value in seconds. Default is 5 seconds. Range is 1 second to indefinitely.
indefinitely	Accumulate link-flap instances forever.
threshold	Sets number of link-flap events tolerated before action is taken.
<i>threshold</i>	Threshold value. Default is 10. Minimum threshold is 1; maximum threshold value depends on the link-flap detection configured link-flap interval and the link scan interval.
disable-time	Sets time period a port remains disabled after detecting excessive link flapping.
<i>disable_time</i>	Disable time in seconds. Default is 300 seconds. Range is 1 seconds to until enabled by user.
until-cleared	Port remains down until you issue <code>clear ports [<i>port_list</i> all] link-flap-detection status</code> command.

Default

These options have the following default values:

Option	Default Value
Interval	5 seconds
Threshold	10 link flaps
Disable time	300 seconds

Usage

If the default link-scan interval is 50 ms, then in 1 second, a maximum of 20 link state transitions (up or down) and 10 link down transitions can be detected. Assuming the link-flap interval is set to 5, the maximum link-flap threshold is $10 * 5 = 50$. Maximum threshold for interval of 10 seconds appears in the output of the `show ports all link-flap configuration` command.

For example, the following sequence of commands generates an error message:

```
configure ports 7 link-flap-detection interval 5
configure ports 7 link-flap-detection threshold 200
Error: Maximum threshold is 100 for port 7 for current
configuration of link-flap interval of 5 seconds.
```

Similarly, if the current threshold is 50, default link-scan interval is 50 ms, and the interval is changed to 4 seconds, then an error message appears:

```
configure ports 7 link-flap-detection threshold 100
configure ports 7 link-flap-detection interval 2
Error: Current threshold of 100 for port 7 is invalid with
new interval value of 2 seconds. Threshold must be less
than 40 for interval to be 2 seconds.
```

Example

The following example sets the threshold value to 15 link flaps that can be accumulated in an infinite interval for all ports.

```
configure ports all link-flap-detection interval indefinitely threshold 15
```

History

This command was first available in ExtremeXOS release 22.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ports link-flap-detection

```
configure ports [port_list | all] link-flap-detection [on | off]
```

Description

Turns on or off link-flap detection.

Syntax Description

ports	Physical ports.
<i>port_list</i>	List of ports to activate link-flap detection upon.
all	Activates link-flap detection on all port in the system.
on	Link-flap detection is on.
off	Link-flap detection is off.

Default

Link-flap detection is disabled by default.

Example

The following example turns off link-flap detection on ports 1–15:

```
Configure ports 1-15 link-flap-detection off
```

History

This command was first available in ExtremeXOS 22.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ports link-scan interval

```
configure ports link-scan interval [ milliseconds | default ] { slot
    [ slot | all ] }
```

Description

Configures the link-scan interval. The configure command allows the user to set the interval in a range between the default for the platform and 500 ms. A higher interval can free up CPU cycles when fast link detection is not a requirement.

Syntax Description

ports	Ports.
link-scan	Configure link scan attributes for polling port status.
interval	Configure amount of time between polling port status
milliseconds	Interval in milliseconds. Range is 50 to 500 for most platforms. The minimum interval depends on the default for the platform."; type="int"; range="[50,500]"
default	Default interval (50 ms for most platforms).
slot	Slot number (default all slots)"; capability="slot_available"
all	All slots.

Default

50 ms.

Usage Guidelines

Use this command to configure the link-scan interval.

Example

```
# sh ports link-scan
Slot   Interval (ms)
-----
  1      50 (default)
  2     300
  3      50 (default)
  4      50 (default)
```

```
5
6
7
8      200
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ports monitor vlan

```
configure ports [port_list|all] monitor vlan [vlan_name | vlan_list]
               {rx-only | tx-only}
```

Description

Starts counting VLAN statistics on a port or a group of ports.

Syntax Description

<i>port_list</i>	Specifies one or more ports. May be in the form: 1, 2, 3-5, 2:5, 2:6-2:8.
all	Specifies all ports.
<i>vlan_name</i>	Specifies a VLAN name.
<i>vlan_list</i>	Specifies a VLAN list of IDs.
rx-only	Specifies receive statistics.
tx-only	Specifies transmit statistics.

Default

N/A.

Usage Guidelines

Use this command to configure access to VLAN statistics per port.

The rx-only and tx-only parameters are intended for, but not restricted to, use on ports that support both receive and transmit statistics. Ports on slots that do not support transmit statistics do not require explicit use of the rx-only keyword. In the absence of specifying either rx-only or tx-only, both RX and TX VLAN statistics are gathered if both are supported on the configured port.

When both receive and transmit statistics are configured and resources for either receive or transmit are not available, neither receive nor transmit statistics will be configured.

The number of VLANs that can be monitored is dependent on filtering resources on the involved switch.

When per-port monitoring is configured, the following commands display the latest statistics directly from the hardware in real time. This information is not logged.

To display VLAN statistics at the port level, use the following command:

```
show ports {port_list} vlan statistics {no-refresh | refresh}
```

To display VLAN statistics at the VLAN level, use the following command:

```
show vlan {vlan_name | vlan_list} statistics
```

Example

The following example configures per-port monitoring of transmit statistics for a set of ports for the VLAN named finance on a switch:

```
configure ports 2,3 monitor vlan finance tx-only
```

History

This command was first available in ExtremeXOS 12.0.

Support for ExtremeSwitching switches was added in ExtremeXOS 12.5.

The `vlan_list` variable was added in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ports partition

```
configure ports [port_list | all] partition [1x100G | 1x40G | 2x50G |
4x10G | 4x25G]
```

Description

Partitions 100G and 40G ports into multiple partition speeds, and partitions 25G ports into a single 10G port for the ExtremeSwitching 5520 platform.

Syntax Description

<i>port_list</i>	Specifies one or more ports.
all	Specifies all ports.
1x100G	Specifies partitioning a 100G port into a single 100G port (applies only to switches with 100G port(s)). This option is not available on the ExtremeSwitching 5520 platform.

1x40G	Specifies partitioning a 40G port into a single 40G port (applies only to switches with 40G port(s)).
2x50G	Specifies partitioning a 100G port into a two 50G port (applies only to switches with 100G port(s)).
3x1G	Specifies partitioning a 5720-VIM-6YE into a three consecutive 1G ports. Apply 3x1G option on the first VIM port (port 27 or 30 for 24-port models and ports 51 or 55 for 48-port models). All three ports (ports 27-29 or 30-32 for 24-port models and ports 51-53 and 54-56 on 48-port models) on the VIM become 1G ports when applying this option.
3x10G	Specifies partitioning a 5720-VIM-6YE into a three consecutive 10G ports. Apply 3x10G option on the first VIM port (port 27 or 30 for 24-port models and ports 51 or 55 for 48-port models). All three ports (ports 27-29 or 30-32 for 24-port models and ports 51-53 and 54-56 on 48-port models) on the VIM become 10G ports when applying this option.
3x25G	Specifies partitioning a 5720-VIM-6YE into a three consecutive 25G ports. Apply 3x25G option on the first VIM port (port 27 or 30 for 24-port models and ports 51 or 55 for 48-port models). All three ports (ports 27-29 or 30-32 for 24-port models and ports 51-53 and 54-56 on 48-port models) on the VIM become 25G ports when applying this option.
4x10G	Specifies partitioning a 40G port into a four 10G port (applies only to switches with 40G port(s)). Specifies partitioning a 25G port into a single 10G port. Apply 4x10G option on the first 5520 VIM port (port 33 for 24-port models and port 57 for 48-port models). All four ports on the 5520 VIM become 10G ports when applying this option.
4x25G	Specifies partitioning a 100G port into a four 25G port. Also specifies changing a 25G port that was partitioned to a 10G port to back to 25G. Apply 4x25G option on the first 5520 VIM port (port 33 for 24-port models and port 57 for 48-port models). All four ports on the 5520 VIM become 25G ports when applying this option.

Default

For 5520-VIM-4YE, 25G ports default to 1x25G.

For 5520 QSFP28 ports the default partition is 40G.

For 5720-VIM-2CE ports default to 1x100G.

For 5720-VIM-6YE ports default to 3x25G.

Usage Guidelines

For the ExtremeSwitching 5520 platform, the QSFP28 ports in Ethernet mode can be configured for 40G (default), 2x50G, 4x25G or 4x10G. These ports are normally configured for stacking. To use these for Ethernet, you must disable stacking support.

For the ExtremeSwitching 5520-VIM-4YE module you can switch all four ports on the VIM from 25G to 10G using the **4x25G** option.

Example

The following example partitions port 6:1 into four 10G ports:

```
# configure ports 6:1 partition 4x10G
```

History

This command was available in ExtremeXOS 12.6.

This command was expanded to include partitioning 100G ports in ExtremeXOS 22.2

Dynamic partitioning (no reboot required) was added in ExtremeXOS 30.2.

Platform Availability

This command is available on ExtremeSwitching 5520 and 5720 switch models:

Table 17: QSFP28 Port Partitioning

Switch model	Port bandwidth	Port partitions
5520 (all models) QSFP28	100 Gb	One QSFP28 port with one of the following: <ul style="list-style-type: none"> Two 50 Gb ports One 40 Gb port (default) Four 25 Gb ports Four 10 Gb ports Note: These ports do not support 1x100Gb
5720 (all models) U1 and U2 ports	100Gb	One U1 or U2 port with one of the following: <ul style="list-style-type: none"> One 40 Gb port One 100 Gb port (default) Two 50 Gb ports Four 10 Gb ports Four 25 Gb ports
5720-VIM-6YE	25G	The 5720-VIM-6YE has two speed groups, each with three ports that can be partitioned with the following: <ul style="list-style-type: none"> Three 1 Gb ports Three 10 Gb ports Three 25 Gb ports (default)
5720-VIM-2CE	100G	The 5720-VIM-2CE has two ports, each of which can be partitioned with the following: <ul style="list-style-type: none"> One 100 Gb port (default) One 40 Gb port Four 10 Gb ports Four 25 Gb ports

configure ports protocol filter

```
configure ports [port_list | all] protocol filter [none | filter_name]
```

Description

Configures protocol filtering on a port.

Syntax Description

<i>port_list</i>	Specifies the port list separated by a comma (,) or dash (-).
all	Specifies all ports.
protocol filter	Specifies the protocol filter.
none	Specifies to not perform protocol filtering on specified ports.
<i>filter_name</i>	Specifies the protocol filter name.

Default

Disabled.

Usage Guidelines

Use this command to configure protocol filtering on a port.

Example

The following example unbinds the L2PT profile from peer 1.1.1.1 of VPLS cust2:

```
configure l2vpn vpls cust2 peer 1.1.1.1 l2pt profile none
```

The following example enables filtering of protocols in *my_list* on port 1:

```
configure ports 1 protocol filter "my_list"
```

The following example disables protocol filtering on port 7:

```
configure ports 7 protocol filter none
```

History

This command was first available in ExtremeXOS 15.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ports qosprofile

```
configure ports port_list {qosprofile} qosprofile
```

Description

Creates a port-based traffic group, which configures one or more ingress ports to use a particular egress QoS profile.

Syntax Description

<i>port_list</i>	Specifies a list of ports or slots and ports.
<i>qosprofile</i>	Specifies a QoS profile.

Default

All ingress ports have the default qosprofile of QP1.

Usage Guidelines

This command assigns traffic ingressing the specified port to a specified egress QoS profile. Extreme switches support eight egress QoS profiles (QP1 to QP8) for each port. SummitStack does not permit configuration of QP7.

Example

The following command configures port 5 to use QoS profile QP3:

```
configure ports 5 qosprofile QP3
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ports rate-limit egress

```
configure ports port_list rate-limit egress [no-limit | cir-rate [Kbps | Mbps | Gbps] {max-burst-size burst-size [Kb | Mb]}]
```

Description

Configures an egress traffic rate limit for a port or groups of ports.

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
no-limit	Specifies traffic be transmitted without limit; use to reconfigure or unconfigure previous rate-limiting parameters.
<i>cir-rate</i>	Specifies the desired rate limit in Kbps, Mbps, or Gbps.
max-burst-size	Specifies the maximum burst size or peak burst size in kilobits (Kb) or megabits (Mb).

Default

No-limit.

Usage Guidelines

Port speed limits the egress traffic, as follows:

- 1 Gbps port—64 Kbps increments.
- 10 Gbps port—1 Mbps increments.

If the specified egress limit (*cir-rate*) is not a multiple of 64 Kbps for a 1 Gbps port or 1 Mbps for a 10Gbps port, the specified value is rounded down to the nearest appropriate multiple based on the port type.

Use the **no-limit** parameter to:

- Unconfigure egress rate limiting on the port(s).
- Reconfigure existing egress rate limiting on the port(s).

The *max-burst-size* parameter is the amount of traffic above the value in the *cir-rate* parameter that is allowed to burst from the port(s) for a short duration. If *max-burst-size* has been configured as "0", then it will use maximum available burst value.

Example

The following command configures egress rate-limiting on port 1 a switch for 3 Mbps and a maximum burst size of 5 M bits:

```
configure port 1 rate-limit egress 3 Mbps max-burst-size 5 Mb
```

History

This command was available in ExtremeXOS 11.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ports rate-limit flood

```
configure ports [port_list | port_group]rate-limit flood [broadcast |
multicast | unknown-destmac] [no-limit | pps {out-actions [{log}
{trap} {disable-port}]}]]
```

Description

Limits the amount of ingress flooded traffic; minimizes network impact of broadcast loops.

Syntax Description

<i>port_list</i>	Specifies the port number. On a stand-alone switch, this value is just the port number, and on a SummitStack, this value is the slot and port number.
<i>port_group</i>	Port group name.
broadcast	Specifies all broadcast packets.
multicast	Specifies all flooded multicast packets (known IP multicast caches are still forwarded at line rate).
unknown-destmac	Specifies all packets with unknown MAC DAs.
no-limit	Specifies unlimited rate.
<i>pps</i>	Packets per second allowed; range is from 0 to 262,144.
out-actions	Out-of-profile action.
log	Generate log event if traffic exceeds configured rate.
trap	Generate <i>SNMP</i> trap if traffic exceeds configured rate.
disable-port	Disable the underlying port when traffic exceeds configured rate.

Default

No limit.

Usage Guidelines

Use this command to limit the amount of ingress flooding traffic and to minimize the network impact of broadcast loops.



Note

When the **multicast** keyword is used, both known and unknown multicast traffic will be rate limited.

To display results, use the `show ports rate-limit flood` command.

Example

The following example rate limits broadcast packets on port 3 on a stand-alone switch to 500 pps:

```
configure ports 3 rate-limit flood broadcast 500
```

History

This command was available in ExtremeXOS 11.1.

The **out-actions**, **log**, **trap**, **disable-port**, and *port_group* options were added in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ports redundant

```
configure ports primaryPort redundant secondaryPort {link [on | off]}
```

Description

Configures a software-controlled redundant port.

Syntax Description

<i>primaryPort</i>	Specifies one primary port or slot and port.
redundantPort <i>secondaryPort</i>	Specifies one or redundant port or slot and port.
link	Specifies state of link: on—Specifies keeping the redundant port active, but block traffic off—Specifies forcing the link down on the redundant port Note: The default value is off.

Default

N/A.

Usage Guidelines

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

A software-controlled redundant port is configured to back up a specified primary port; both ports are on the same device. 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 configure the redundant link to be always physically up but logically blocked or to be always physically down. The default is off, or the redundant link is down.

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

- You can configure only one redundant port for each primary port.
- You cannot have any Layer 2 protocols configured on any of the VLANs that are present on the ports. (You will see an error message if you attempt to configure software redundant ports on ports with VLANs running Layer 2 protocols.)
- The primary and redundant port must have identical VLAN memberships.
- The master port is the only port of a load-sharing group that can be configured as either a primary or redundant port. (The entire trunk must go down before the software-controlled redundant port takes effect.)
- Only one side of the link should be configured as redundant.

Example

The following command configures a software-controlled redundant port:

```
configure ports 1:3 redundant 2:3
```

History

This command was available in ExtremeXOS 11.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ports vlan

```
configure ports port_list [ {tagged tag} vlan vlan_name | {tagged}  
      vlan vlan_list ] [limit-learning number {action [blackhole | stop-learning] } | lock-learning | unlimited-learning | unlock-learning]
```

Description

Configures virtual ports for limited or locked MAC address learning.

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
tagged <i>tag</i>	Specifies the port-specific <u>VLAN</u> tag. When there are multiple ports specified in the <i>port_list</i> , the same tag is used for all of them.
<i>vlan_name</i>	Specifies the name of the VLAN.
<i>vlan_list</i>	Specifies a VLAN list of IDs.

limit-learning <i>number</i>	Specifies a limit on the number of MAC addresses that can be dynamically learned on the specified ports.
blackhole	Specifies that blackhole entries are created for MAC addresses that exceed the limit-learning limit. This is the default setting.
stop-learning	Specifies that the learning be halted to protect the switch from exhausting <i>FDB</i> resources by not creating blackhole entries.
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

If you have enabled *ESRP*, see the appropriate volume of the [Switch Engine 32.3 User Guide](#) for information about using this feature with ESRP.

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 *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 fdb` and `delete fdb` commands. These override any dynamically learned entries.

For ports that have a learning limit in place, the following traffic still flows 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 participate in an Extreme Standby Router Protocol (ESRP) domain, 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 protocol data units (PDUs) from being dropped due to MAC address limit settings.

Stop learning

When stop-learning is enabled with learning-limit configured, the switch is protected from exhausting FDB resources by not creating blackhole entries. Any additional learning and forwarding is prevented, but packet forwarding from FDB entries is not impacted.

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 still flows 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 display the locked entries on the switch, use the following command:

```
show fdb
```

Locked MAC address entries have the "I" flag.

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

```
show vlan vlan name security show ports port_list info detail
```

Example

The following example 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:

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

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

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

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

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

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

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

The following example illustrates use of the **tagged** keyword:

```
configure ports 1 tag 10 vlan accounting learning-limit 128
configure ports 1 vlan accounting learning-limit 128
configure ports 4 tag 10 vlan accounting lock-learning
configure ports 4 vlan accounting lock-learning
```

History

This command was first available in ExtremeXOS 11.1.

The `vlan_list` option was added in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure power monitor

```
configure power monitor poll-interval [off | seconds] change-action
                        [none | [log | log-and-trap | trap] change-threshold watts]
```

Description

Configures the power visualization, which periodically polls for input power usage.

Syntax Description

seconds	Input power usage poll interval in seconds. If zero is configured, then the input power measurement is disabled.
change-action	The action to be taken whenever the power is increased or decreased by the configured threshold power value (none, log, log-and-trap, or trap).
watts	The power value in watts for the change threshold. The default value is 2 watts.

Default

The default poll interval is 60 seconds.

The default change action is none.

The default change threshold is 2 watts.

Usage Guidelines

Use this command to configure change actions to be taken when input power usage is increased or decreased by the configured threshold power value. The polling interval is also configurable, with a default value of 60 seconds.



Note

Input power usage values are only estimates.

Example

The following command configures a polling interval of 10 seconds, a change action of log-and-trap, and a change threshold of 3 watts:

```
configure power monitor poll-interval 10 change-action log-and-trap change-threshold 3
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure private-vlan add network

```
configure private-vlan name add network vlan_name
```

Description

Adds the specified VLAN as the network VLAN on the specified PVLAN.

Syntax Description

<i>name</i>	Specifies the name of the PVLAN to which the VLAN is added.
<i>vlan_name</i>	Specifies a VLAN to add to the PVLAN.

Default

N/A.

Usage Guidelines

The VLAN must be created and configured with a tag before it is added to the PVLAN.

Example

The following example adds VLAN "sharednet" as the network VLAN for the PVLAN named "companyx":

```
configure private-vlan companyx add network sharednet
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms that support the Private VLAN feature. The features and the platforms that support them are listed in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure private-vlan add subscriber

```
configure private-vlan name add subscriber vlan_name {non-isolated}  
                        {loopback-port port}
```

Description

Adds the specified VLAN as a subscriber VLAN on the specified PVLAN.

Syntax Description

<i>name</i>	Specifies the name of the PVLAN to which the VLAN is added.
<i>vlan_name</i>	Specifies a VLAN to add to the PVLAN.
non-isolated	Configures the subscriber VLAN as a non-isolated subscriber VLAN.
<i>port</i>	Specifies the port that serves as the loopback port.

Default

If the **non-isolated** option is omitted, this command adds the specified VLAN as an isolated subscriber VLAN.

Usage Guidelines

The VLAN must be created and configured with a tag before it is added to the PVLAN. If the non-isolated option is omitted, the VLAN is added as an isolated subscriber VLAN. If the non-isolated option is included, the VLAN is added as a non-isolated subscriber VLAN.

If two or more subscriber VLANs have overlapping ports (where the same ports are assigned to both VLANs), each of the subscriber VLANs with overlapping ports must have a dedicated loopback port.

Example

The following example adds VLAN "restricted" as a subscriber VLAN for the PVLAN named "companyx":

```
configure private-vlan companyx add subscriber restricted isolated
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms that support the Private VLAN feature. For features and the platforms that support them, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure private-vlan delete

```
configure private-vlan name delete [network | subscriber] vlan_name
```

Description

Deletes the specified VLAN from the specified PVLAN.

Syntax Description

<i>name</i>	Specifies the name of the PVLAN from which the VLAN is deleted.
network	Specifies that the VLAN to be deleted is a network VLAN.
subscriber	Specifies that the VLAN to be deleted is a subscriber VLAN.
<i>vlan_name</i>	Specifies the VLAN to delete from the PVLAN.

Default

N/A.

Usage Guidelines

This command deletes a VLAN from a PVLAN, but it does not delete the VLAN from the system—it just breaks the link between the VLAN and the PVLAN. You can use this command to delete both network and subscriber VLANs.

Example

The following example deletes network VLAN "sharednet "from the PVLAN named "companyx":

```
configure private-vlan companyx delete network sharednet
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms that support the Private VLAN feature. For features and the platforms that support them, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure protocol add

```
configure protocol {filter} filter_name add [etype | llc | snap] hex
    {[etype | llc | snap] hex}
```

Description

Configures a user-defined protocol filter.

Syntax Description

filter	Configures a protocol filter.
<i>filter_name</i>	Specifies a protocol filter name.
add	Specifies that you add a protocol.
delete	Specifies that you delete a protocol.
etype	Specifies an ethernet protocol.
llc	Specifies LLC protocol.
snap	Specifies SNAP protocol.
<i>hex</i>	Specifies a four-digit hexadecimal number between 0 and FFFF that represents: <ul style="list-style-type: none"> 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

Supported protocol types include:

- etype—IEEE Ethertype.
- llc—LLC Service Advertising Protocol.
- snap—Ethertype inside an IEEE SNAP packet encapsulation.

A maximum of 16 customized protocol filters can be active at a time.

The protocol filter must already exist before you can use this command. Use the `create protocol` command to create the protocol filter.



Note

Protocol-based VLAN for Etype from 0x0000 to 0x05ff are not classifying as per filter. When traffic arrive with these Etypes, it is classified to native VLAN rather protocol-based VLAN.

Example

The following example adds MPLS to "my_filter":

```
configure protocol "my_filter" add etype 0x8847
configure protocol filter "my_filter" add etype 0x8847
```

The following example deletes MPLS from "my_other_filter":

```
configure protocol "my_other_filter" delete etype 0x8847
configure protocol filter "my_other_filter" delete etype 0x8847
```

History

This command was first available in ExtremeXOS 10.1.

The **filter** keyword and options were added in ExtremeXOS 15.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure process group other cpu-limit

```
configure process group other cpu-limit cpu_limit
```

Description

This command changes the CPU limit for the "Other" (non-vital) process group.

Syntax Description

other	Designates "Other" (non-vital) process group.
cpu-limit	Designates changing the maximum amount of CPU that the "Other" process group can use during resource contention.
<i>cpu-limit</i>	Sets the value for the CPU limit value as a percentage. The valid range is 5% to 50%; default is 10%.

Default

By default, the CPU limit of "Other" group is 10%. With the default configuration, the "Vital" group CPU limit is 90%.

Usage Guidelines

This command allows you to configure CPU limits for the "Other" group. The configured CPU percentage is guaranteed for the "Other" group, unless a real-time kernel task needs CPU.

When this command is issued, the CPU limit for the "Vital" group is changed as well. For example, if you change the CPU limit value to 30, the new values are: 70% for "Vital", and 30% for "Other".

If you try to configure a limit that is greater than the current configured value, a warning message appears:

Warning: Increasing CPU limit of the "Other" group may degrade EXOS performance and lead to network instability. The CPU limit for the "Other" group has been increased from 10% to 30%.

To see the status of the process groups, use the command [show process group](#) on page 3137.

Example

The following example changes the "Other" process group CPU limit to 30%. Additionally, the "Vital" group is changed to 70%:

```
# configure process group other cpu-limit 30
```

History

This command was first available in ExtremeXOS 22.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure process group other memory-limit

```
configure process group other memory-limit memory_limit
```


Description

This command changes the memory limit for the "Other" (non-vital) process group.

Syntax Description

other	Designates "Other" (non-vital) process group.
memory-limit	Designates changing the memory limit for the "Other" process group.
<i>memory-limit</i>	Sets the value, as a percentage, for the "Other" process group memory limit. The valid range is 5% to 50%; default is 5% of total system memory.

Default

Default memory limit for the "Other" group is 5% of total system memory. With the default configuration, the memory limit of the "Vital" group is 95%.

Usage Guidelines

This command allows you to increase or decrease the memory limit assigned to the "Other" (non-vital) process group. The configured limit is used as the new upper bound for the "Other" group. When this command is issued, the memory limit for the "Vital" group is changed as well. For example, if the current value is 95% for "software application", and 5% for "Other", if you change the memory limit value to 30, the new values are: 70% for "Vital", and 30% for "Other".

When you issue this command, a warning message appears:

```
Warning: Increasing memory-limit of the "Other" group will reduce the available memory for "EXOS".
```

If you try to set a memory limit below the value that is already consumed by the "Other" group, an error message appears. For example, when you change the memory limit to 5% when it is already consuming 8.7%, the following error message appears.

```
Error: Desired memory-limit (5%) must be greater than or equal to the current memory consumption (8.7%) of the group "Other".
```

You also cannot increase the memory limit on a process group beyond the available memory for the process group. For example, if you try increasing the memory limit on the "Other" (non-vital) group to 40% when the group is already consuming 70%, the following error message appears:

```
Error: Desired memory-limit (40%) must be less than or equal to the available memory of (30%) for the "Other" group. "EXOS" is currently consuming 70% of system memory.
```

To see the status of the process groups, use the command [show process group](#) on page 3137.

Example

The following example sets the "other" process group memory limit to 25%. This also sets the memory limit to 75% for the "Vital" group:

```
# configure process group other memory-limit 25
```

History

This command was first available in ExtremeXOS 22.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure protocol delete

```
configure protocol name delete [etype | llc | snap] hex { [etype | llc | snap] hex } ...
```

Description

Deletes the specified protocol type from a protocol filter.

Syntax Description

<i>name</i>	Specifies a protocol filter name.
<i>hex</i>	Specifies a four-digit hexadecimal number between 0 and FFFF that represents: <ul style="list-style-type: none"> 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

Supported protocol types include:

- etype—IEEE Ethertype.
- llc—LLC Service Advertising Protocol.
- snap—Ethertype inside an IEEE SNAP packet encapsulation.

Example

The following example deletes protocol type LLC SAP with a value of FEFF from protocol "fred":

```
configure protocol fred delete llc feff
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure protocol filter

```
configure protocol filter filter_name [add | delete] dest-mac
    mac_address {[etype | llc | snap] hex} {field offset offset value
    value {mask mask}} {tagged}
```

Description

Configures the destination address as well as an arbitrary field of the protocol.

Syntax Description

<i>filter_name</i>	Specifies a protocol filter name.
add	Specifies that you add a protocol.
delete	Specifies that you delete a protocol.
dest-mac	Specifies the destination MAC address used by PDUs of the protocol.
<i>mac_address</i>	Specifies the MAC address.
etype	Specifies the EtherType used by PDUs of the protocol.
llc	Specifies the LLC DSAP and SSAP used by PDUs of the protocol.
snap	Specifies the SNAP protocol identifier used by PDUs of the protocol.
<i>hex</i>	Specifies a four-digit hexadecimal number between 0 and FFFF that represents: <ul style="list-style-type: none"> 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.
field	Specifies a field used by PDUs of the protocol.
offset	Specifies the offset of the field from the start of the PDU.
value <i>value</i>	Specifies the value of the field in hexadecimal (for example, A1:B2:0C. Maximum 16 bytes).
mask <i>mask</i>	Specifies the mask for the field in hexadecimal (for example, FF:FF:0F. Maximum 16 bytes).
tagged	Specifies if the protocol is a tagged protocol. Default is not tagged.

Default

N/A.

Usage Guidelines

Supported protocol types include:

- etype—IEEE Ethertype.
- llc—LLC Service Advertising Protocol.
- snap—Ethertype inside an IEEE SNAP packet encapsulation.

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.

No more than seven protocols can be active and configured for use.



Note

Protocol-based VLAN for Etype from 0x0000 to 0x05ff are not classifying as per filter. When traffic arrive with these Etypes, it is classified to native VLAN rather than protocol-based VLAN.

Example

The following example LACP to the protocol list "mylist":

```
configure protocol "mylist" add dest-mac 01:80:C2:00:00:02 etype 0x8809 field offset 14
value
01 mask FF
```

The following example removes EFM OAM from the protocol list "mylist":

```
configure protocol filter "mylist" delete dest-mac 01:80:C2:00:00:02 etype 0x8809 field
offset
14 value 03 mask FF
```

The following example configures a mismatched mask and value:

```
configure protocol "mylist" delete dest-mac 01:80:C2:00:00:02 etype 0x8809 field offset
14 value 03 mask FF:FF
Error: The length of the field value is not the same as the field mask.
```

History

This command was first available in ExtremeXOS 15.5.

The **tagged** keyword was added in ExtremeXOS 31.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure qosprofile

```
configure qosprofile egress qosprofile [{minbw minbw_number} {maxbw
maxbw_number} | {peak_rate peak_bps [K | M]}] [ports [port_list |
port_group |all]]

configure qosprofile qosprofile [{minbw minbw_number} {maxbw
maxbw_number} | {{committed_rate committed_bps [K | M]} {peak_rate
peak_bps [K | M]} | [ports [port_list | all]]]

configure {qosprofile} qosprofile [{maxbuffer buffer_percentage} {weight
weight_value | use-strict-priority} {ports [port_list | port_group |
all]}]
```

Description

Modifies the default egress [QoS](#) profile parameters.

Syntax Description

minbw	The minimum bandwidth (minbw) option specifies the committed information rate as a percentage of the maximum port speed. The range is 0 to 100%, and the default value is 0. When autonegotiation is off, the CIR is the specified percentage of the configured port speed. When autonegotiation is on, the CIR is the specified percentage of the maximum port speed.
maxbw	The maximum bandwidth (maxbw) option specifies the peak rate as a percentage of the maximum port speed. The range is 0 to 100%, and the default value is 100. When autonegotiation is off, the peak rate is the specified percentage of the configured port speed. When autonegotiation is on, the peak rate is the specified percentage of the maximum port speed (the switch does not detect the negotiated port speed).
peak_rate	Specifies a peak rate in Kbps (k) bits or Mbps (m).
committed_rate	Specifies a committed information rate in Kbps (k) bits or Mbps (m).
<i>port_list</i>	Specifies a list of slots and ports to which the parameters apply. Specify ports in the following formats: 3-5, 2:5, 2:6-2:8.
<i>buffer_percentage</i>	When used without a port-list, specifies the percentage of the total buffer you are reserving for this QoS profile on all ports for which an override has not been configured. The range is 1 to 100; the default setting is 100. When used with a port-list, specifies a percentage override of the maxbuffer setting for the QoS profile specified. The range is 1-10000; the default is 100 (i.e., no override). Setting 100% is equivalent to unconfiguring the maxbuffer override.
qosprofile	Specifies a QoS profile name.
use-strict-priority	When the global qosscheduler configuration (configure qosscheduler command) is set to weighted-round-robin, this option overrides the global configuration for the specified QoS profile, so that it operates in strict-priority-mode. This enables hybrid strict-priority and weighted-round-robin scheduling operation.

<i>weight-value</i>	Specifies the weight value used for queue service weighting in the weighted-round-robin scheduler for this QoS profile. Range is 1-15 or 1-127 depending on hardware type. 0=strict-priority. Default is 1. This command enables the user to input a weight for queues in the weighted-round-robin scheduler or weighted-deficit-round-robin scheduler. The weight of both WRR and WDRR algorithms have been extended to 1-127.
ports	Port list for maxbuffer and per-port weight override.
<i>port_list</i>	Port list.
<i>port_group</i>	Port group name.
all	Specifies this applies to all ports on the device.

Default

- QoS profiles—QP1 and QP8 on SummitStack and ExtremeXOS series switches
- Minimum bandwidth—0%
- Maximum bandwidth—100%
- Maximum buffer—100%
- Maxbuffer override—100% (no override)
- Weight—1
- Priority—By default, each qosprofile is assigned a different priority level:
 - QP1 - 1, Low (the lowest priority)
 - QP2 - 2, LowHi
 - QP3 - 3, Normal
 - QP4 - 4, NormalHi
 - QP5 - 5, Medium
 - QP6 - 6, MediumHi
 - QP7 - 7, High
 - QP8 - 8, HighHi (highest priority)

Usage Guidelines



Note

You can view the effect of setting the buffer-percentage using the `show ports port-list buffer` command.



Note

You can view the configured buffer-percentage value using the `show qosprofile` or `show qosprofile ports port-list` commands, respectively.

The maximum bandwidth value can be configured as either:

- An absolute percentage of the total maximum link speed, regardless of the currently configured or negotiated speed, OR
- An absolute peak rate in Mbps or Kbps.

QoS profiles QP1 and QP8 are preconfigured. If you want to use a QoS profile in the range of QP2 through QP7, you must first create the QoS profile. QoS profile QP7 is reserved on SummitStack for stack management and cannot be created or modified.

When specified without a port-list, the `maxbuffer` parameter can configure a reduction in the maximum amount of packet buffer space allotted to the specified QoS profile. If you reduce the allotment below the default value of 100%, the reduction releases packet buffer space to the shared packet buffer. Regardless of the setting for this parameter, the system does not drop any packets as long as reserved packet buffer memory for the port and QoS profile or shared packet memory for the port remains available.

**Note**

The configuration defined by the `maxbuffer` attribute in this command can be overridden on a per-port basis if the port is specified along with the `maxbuffer` parameter.

When specified with a port-list, the `maxbuffer` setting overrides the system-wide reduction of packet buffer reservation set with the `configure qosprofile maxbuffer` command for the specified QoS profile. If the packet buffer reservation is reduced to 75 percent for the entire QoS profile, the specified ports are allotted 75% of the allotment for the specified QoS profile. If for specified ports the `maxbuffer` is set to 200 percent, the packet buffer reservation will be set to 200 percent of the normal packet buffer reservation for those ports, thus overriding the `maxbuffer` percentage set for the QoS profile.

**Note**

The packet buffer configuration feature is provided for expert users who fully understand the impact of buffer configuration changes. Improper buffer configuration can stop traffic flow through QoS profiles and ports for which no direct configuration change was made.

A range of ports has its own packet buffer pool. The `maxbuffer` override capability allows you to overcommit the packet buffer pool for the port range. When a packet buffer pool is overcommitted by more than 20%, the following message appears in the system log:

```
Warning: Packet memory is overcommitted by <percentage> for ports in range <port-range>
```

It is also possible to configure `maxbuffer` overrides such that the size of the shared portion of the buffer pool is reduced to zero. If some port and QoS profile in the port range for that buffer pool does not have sufficient reserved packet memory to accommodate larger packets, it will be impossible for that port and QoS profile to transmit any packets of the larger size. In this case, the following message appears in the system log:

```
Warning: At least one port and QoS profile in port range <port-range> cannot transmit packets larger than <packet-size> because of packet memory configuration.
```

The `weight-value` parameter does not apply when the switch is configured for strict priority scheduling, which is the default configuration. To configure the type of scheduling you want to use for the entire switch, use the `configure gosscheduler` command.

The `weight-value` parameter configures the relative weighting for each QoS profile. Because each QoS profile has a default weight of 1, all QoS profiles have equal weighting. If you configure a QoS profile with a weight of 4, that specified QoS profile is serviced 4 times as frequently as the remaining QoS profiles, which still have a weight of 1. If you configure all QoS profiles with a weight of 16, each QoS profile is serviced equally but for a longer period.

When the switch is configured for weighted-round-robin mode, the use-strict-priority option overrides the switch configuration for the specified QoS profile on all ports. Among QoS profiles configured with the use-strict-priority-option, QoS profile QP8 has the highest priority and QP1 has the lowest priority. All strict-priority QoS profiles are serviced first according to their priority level, and then all other QoS profiles are serviced based on their configured weight.



Note

If you specify use-strict-priority, lower-priority queues and weighted-round-robin queues are not serviced at all as long as higher-priority queues have any remaining packets.

Example

The following example overrides the maximum buffer setting configured on QoS profile qp1 for port1:1:

```
# configure qosprofile qp1 maxbuffer 75 port 1:1
```

History

This command was first available in ExtremeXOS 10.1.

Committed and peak rates were added in ExtremeXOS 11.0. Also in ExtremeXOS 11.0, ports were made mandatory.

Support for all platforms was added in the respective platform introduction releases.

The use-strict-priority option was added in ExtremeXOS 12.3.

The ability to configure a maxbuffer override was added in ExtremeXOS 12.5.

The `port_group` variable was added in ExtremeXOS 16.1.

Platform Availability

This command is available on all platforms with specific parameter exceptions as noted in the Syntax Description above.

configure qosprofile weight

```
configure qosprofile qp8 weight weight_value
```

Description

This command enables the user to input a weight value for queue service weighting in the weighted-round-robin scheduler or weighted-deficit-round-robin scheduler for this QoS profile. The weight value of both WRR and WDRR algorithms have been extended to 1-127 on this supported hardware (refer to the [Switch Engine 32.3 User Guide](#) for supported hardware).

Syntax Description

<code>weight_value</code>	Range is 1-15 or 1-127 depending on hardware type.
---------------------------	--

Default

Strict priority.

Usage Guidelines

Use this command to input a weight value for queue service weighting in the weighted-round-robin scheduler or weighted-deficit-round-robin scheduler for this QoS profile. The weight value of both WRR and WDRR algorithms have been extended to 1-127 on this supported hardware (refer to the [ExtremeXOS 22.6 User Guide](#) for supported hardware).

Example

Example output not yet available and will be provided in a future release.

History

This command was first available in ExtremeXOS 12.7.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure qosprofile wred

```
configure {qosprofile} {egress} qosprofile [wred [{color [tcp [green
| red] | non-tcp [any|red]] [{min-threshold min_thresh} {max-
threshold } {max-drop-rate max_drop_rate}}] | avg-weight avg_weight]]
ports [port_list |all]
```

Description

Configures WRED on the specified [QoS](#) profile for the specified port.

Syntax Description

egress	This optional parameter specifies an egress QoS profile.
<i>qosprofile</i>	Specifies a QoS profile name. Valid names are QP1 to QP8.
color	Specifies the WRED color to be configured.
green	Specifies that the WRED configuration applies to TCP traffic that is marked green.
non-tcp any	Specifies that the WRED configuration applies to any non-TCP traffic.
red	Specifies that the WRED configuration applies to TCP traffic that is marked red.
<i>min_thresh</i>	Specifies the minimum threshold for the specified WRED color. The range is 1 to 100 percent.

max_threshold	Specifies the maximum threshold for the specified WRED color. The range is 1 to 100 percent.
<i>max_drop_rate</i>	Specifies the maximum drop rate for the specified WRED color. The range is 1 to 100 percent.
<i>port_list</i>	Specifies a list of slots and ports to which the parameters apply. Specify ports in the following formats: 3-5, 2:5, 2:6-2:8.
non-tcp red	Specifies that the WRED configuration applies to non-TCP traffic that is marked red.
<i>avg_weight</i>	Specifies the weight constant for calculating the average queue size for the specified QoS profile. The range is 1 to 15.
all	Specifies that this command applies to all ports on the device.

Default

- Minimum threshold—100%
- Maximum threshold—100%
- Maximum drop rate—100%
- Average weight—4

Usage Guidelines

The *max_drop_rate*, *min_threshold*, and *max_threshold* parameters apply to the specified color. The **avg_weight** parameter applies to all colors on the specified QoS profile. Increasing the *avg_weight* value reduces the probability that traffic is dropped. Conversely, decreasing the *avg_weight* value increases the probability that traffic is dropped.

Example

The following example configures WRED settings for port 2:1, QoS profile qp3, color green:

```
configure qosprofile qp3 wred color tcp green min-threshold 80 max-threshold 95 max-drop-rate 75 ports 2:1
```

The following example configures the average weight for port 2:1, QoS profile qp2:

```
configure qosprofile qp2 wred avg-weight 4 ports 2:1
```

The following example configures WRED settings for non-TCP traffic on port 4, QoS profile qp3:

```
configure qosprofile qp3 wred color non-tcp any min-threshold 10 ports 4
```

The following example configures WRED settings using "wredGroup" as the *port_group* variable:

```
configure qosprofile qp8 wred color tcp red min-threshold 25 max-streshold 75 max-drop-rate 30 ports wredGroup
```

History

This command was first available in ExtremeXOS 12.7.

The `port_group` variable was added in ExtremeXOS 16.1.

Platform Availability

This command is available on the ExtremeSwitching 5520 platform.

configure qosprofile egress wred ecn

```
configure qosprofile egress qp_num wred ecn [on | off] ports [port_list
| all]
```

Description

This command turns Explicit Congestion Notification (ECN) on or off for the corresponding QoS profile for the given port(s).

Syntax Description

<i>qp_num</i>	QoS profile (qp1, qp2, qp3, qp3, qp4, qp5, qp6, qp7, qp8)
wred	Designates weighted random early detection (WRED).
ecn	Designates ECN.
on	Enables ECN.
off	Disables ECN.
ports	Selects ports.
<i>port_list</i>	Selects specific ports to apply the ECN setting for the designated QoS profile.
all	Selects all ports to apply the ECN setting for the designated QoS profile.

Default

N/A.

Usage Guidelines

Weighted Random Early Detection (WRED) drops the packets, based on the average length exceeding a specific threshold value to indicate congestion. Explicit Congestion Notification (ECN) is an extension to WRED that marks the drop-eligible packets, instead of dropping, using the same criteria of minimum threshold, maximum threshold, and drop probability

Example

The following example enables ECN for QoS profile 5 on port 2:

```
# configure qosprofile egress qp5 wred ecn on ports 2
```

History

This command was first available in ExtremeXOS 22.3.

Platform Availability

This command is available on the ExtremeSwitching 5520 platform.

configure qosscheduler weighted-deficit-round-robin

```
configure qosscheduler [strict-priority | weighted-round-robin |
weighted-deficit-round-robin ] {ports [port_list | port_group |
all ] }
```

Description

This command specifies the scheduling algorithm that the switch uses to service [QoS](#) profiles.

Syntax Description

strict-priority	Specifies the switch services the higher-priority QoS profiles first.
weighted-round-robin	Specifies the switch services all QoS profiles based on the configured weighting for each QoS profile.
weighted-deficit-round-robin	Allows you to use a credit-based algorithm in order to sample the size of the packet while scheduling various queues.
ports	Ports to display.
<i>port_list</i>	Port list.
<i>port_group</i>	Port group name
all	All ports.

Default

Strict-priority.

Usage Guidelines

When issued without a *port_list* or *port_group*, this command configures the global scheduling algorithm that will be applied to all ports that have not been configured with per-port scheduling. When issued with a *port_list* or *port_group*, this command configures the scheduling algorithm for specific ports.

The scheduling algorithm for a qosprofile can be overridden either globally or on a per-port basis with the command:

```
configure qosprofile qosprofile use-strict-priority
```

In strict-priority mode, QoS profile QP8 has the highest priority and QP1 has the lowest priority.



Note

Queues are serviced using the configured scheduling algorithm until all of the minBws are satisfied, then all queues are serviced using the configured scheduling algorithm until all of the maxBws are satisfied.

Example

The following example configures the switch for weighted-round-robin servicing:

```
configure qosscheduler weighted-round-robin
```

The following example configures the switch for weighted-deficit-round-robin servicing:

```
configure qosscheduler weighted-deficit-round-robin
```

This command specifies the scheduling algorithm the switch uses to service QoS profiles. Weighted-deficit-round-robin mode of scheduling allows you to use a credit based algorithm in order to sample in the size of the packet while scheduling various queues.

History

This command was first available in ExtremeXOS 15.1.

The **ports** and **all** keywords, and *port_list* and *port_group* variables were added in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure radius algorithm

```
configure radius algorithm [standard | round-robin]
```

Description

This command is used to configure the algorithm used to determine the rotation of *RADIUS* servers.

Syntax Description

standard	Standard Extreme retransmission algorithm.
round-robin	Simple Round Robin retransmission algorithm.

Default

Standard.

Usage Guidelines

Use this command to configure the algorithm to determine rotation of RADIUS servers.

History

This command was first available in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure radius retries

```
configure radius {mgmt-access [primary | secondary] | netlogin [primary
| secondary] | index} retries retries
```

Description

This command is used to set the number of retries the switch will attempt. This value may be global or on a per server basis.

Syntax Description

mgmt-access	<i>RADIUS</i> authentication for management access.
netlogin	RADIUS authentication for netlogin access.
primary	Primary server.
secondary	Secondary server.
<i>index</i>	RADIUS server index.
retries	RADIUS server retries.
<i>retries</i>	RADIUS sever retries. Range 1-20.

Default

The default value is 3, with a range of 0-10.

Usage Guidelines

None.

History

This command was first available in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure radius server client-ip

```
configure radius {mgmt-access | netlogin} [primary | secondary | index]
  server [host_ipaddr | host_ipV6addr | hostname] {udp_port | tls
    {tls_port}} client-ip [client_ipaddr | client_ipV6addr] {vr vr_name}
    {shared-secret {encrypted} secret}
```

Description

This command configures up to eight *RADIUS* authentication servers.



Note

It is recommended to enable loopback mode on the *VLAN* associated with radius if the radius connectivity is established via a front panel port on a SummitStack.

Syntax Description

mgmt-access	Specifies the RADIUS authentication server for switch management.
netlogin	Specifies the RADIUS authentication server for network login.
primary	Configures the primary RADIUS authentication server.
secondary	Configures the secondary RADIUS authentication server.
<i>index</i>	RADIUS server index. Range: 1 - 2147483641.
<i>ipaddress</i>	The IP address of the server being configured.
<i>host_ipV6addr</i>	Server IPv6 address.
<i>hostname</i>	The host name of the server being configured.
<i>udp_port</i>	The UDP port to use to contact the RADIUS authentication server.
tls	Specifies using Transfer Layer Security (TLS).
<i>tls_port</i>	The TLS port to use to contact the RADIUS authentication server.
<i>ipaddress</i>	The IP address used by the switch to identify itself when communicating with the RADIUS authentication server.
<i>client_ipV6addr</i>	Client IPv6 address.
<i>vr_name</i>	Specifies the virtual router on which the client IP is located. Note: User-created VRs are supported only on the platforms listed for this feature in the Switch Engine 32.3 Feature License Requirements document.
shared-secret	Shared secret

<i>secret</i>	Secret string. Important: Use quotes to enclose the string. Failure to do so causes the CLI to treat the string as a comment, since the string starts with a "#" symbol.
encrypted	Password is encrypted.

Default

The following lists the default behavior of this command:

- The UDP port setting is 1812.
- The TLS port setting is 2083.
- The virtual router used is VR-Mgmt, the management virtual router.
- Switch management and network login use the same primary and secondary RADIUS servers for authentication (only if the realm is not specified in the command),.

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.

Beginning with ExtremeXOS 11.2, you can specify one pair of RADIUS authentication servers for switch management and another pair for network login. To specify RADIUS authentication servers for switch management (Telnet, SSH, and console sessions), use the **mgmt-access** keyword. To specify RADIUS authentication servers for network login, use the **netlogin** keyword. If you do not specify a keyword, switch management and network login use the same pair of RADIUS authentication servers.

If you are running ExtremeXOS 11.1 or earlier and upgrade to ExtremeXOS 11.2, you do not lose your existing RADIUS server configuration. Both switch management and network login use the RADIUS authentication server specified in the older configuration.

Specifying **mgmt-access** or **netlogin** before the index will create a RADIUS entry with only that realm specified, if neither are specified both realms will be enabled.



Note

You cannot use a stacking alternate IP address as the RADIUS client in primary RADIUS server configuration.

Example

The following example configures the primary RADIUS server on host radius1 using the default UDP port (1812) for use by the RADIUS client on switch 10.10.20.30 using a virtual router interface of VR-Default:

```
configure radius primary server radius1 client-ip 10.10.20.30 vr vr-Default
```


The following example configures the primary RADIUS server for network login authentication on host netlog1 using the default UDP port for use by the RADIUS client on switch 10.10.20.31 using, by default, the management virtual router interface:

```
configure radius netlogin primary server netlog1 client-ip 10.10.20.31
```

History

This command was first available in ExtremeXOS 10.1.

The **mgmt-access** and **netlogin** keywords were added in ExtremeXOS 11.2.

The *index*, *host_ipv6addr*, *client_ipv6addr*, **shared-secret**, and **encrypted** keywords were added in ExtremeXOS 16.1.

The **tls** keyword with *tls_port* variable was added in ExtremeXOS 31.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure radius shared-secret

```
configure radius [primary | secondary index] shared-secret
                {encryptedencrypted_secret | secret}
```

Description

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

Syntax Description

mgmt-access	Specifies the switch management RADIUS authentication server.
netlogin	Specifies the network login RADIUS authentication server.
primary	Configures the authentication string for the primary RADIUS server.
secondary	Configures the authentication string for the secondary RADIUS server.
<i>index</i>	RADIUS server index.
encrypted	Indicates that the string is already encrypted.
<i>secret</i>	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.

The **mgmt-access** keyword specifies the RADIUS server used for switch management authentication.

The **netlogin** keyword specifies the RADIUS server used for network login authentication.

If you do not specify the **mgmt-access** or **netlogin** keywords, the secret applies to both the primary or secondary switch management and netlogin RADIUS servers.

The **encrypted** keyword is primarily for the output of the show configuration command, so the shared secret is not revealed in the command output. Do not use it to set the shared secret.

Example

The following example configures the shared secret as "purplegreen" on the primary RADIUS server for both switch management and network login:

```
configure radius primary shared-secret purplegreen
```

The following example configures the shared secret as "redblue" on the primary switch management RADIUS server:

```
configure radius mgmt-access primary shared-secret redblue
```

History

This command was first available in ExtremeXOS 10.1.

The **encrypted** keyword was added in ExtremeXOS 11.0.

The **mgmt-access** and **netlogin** keywords were added in ExtremeXOS 11.2.

The *index* variable was added in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure radius timeout

```
configure radius {mgmt-access {primary | secondary} | netlogin {primary | secondary} | index } timeout sec
```

Description

Configures the timeout interval for *RADIUS* authentication requests.

Syntax Description

mgmt-access	Specifies the switch management RADIUS authentication server.
netlogin	Specifies the network login RADIUS authentication server.
primary	Primary server.
secondary	Secondary server.
<i>index</i>	RADIUS server index.
<i>seconds</i>	Specifies the number of seconds for authentication requests. Range is 1 to 240 seconds.

Default

The default is 3 seconds.

Usage Guidelines

This command configures the timeout interval for RADIUS authentication requests. When the timeout has expired, another authentication attempt will be made. After three failed attempts to authenticate, the alternate server will be used. This only refers to the default configuration. After six failed attempts, local user authentication will be used.

The **mgmt-access** keyword specifies the RADIUS server used for switch management authentication.

The **netlogin** keyword specifies the RADIUS server used for network login authentication.

If you do not specify the **mgmt-access** or **netlogin** keywords, the timeout interval applies to both switch management and **netlogin** RADIUS servers.

Example

The following example configures the timeout interval for RADIUS authentication to 10 seconds. After 30 seconds (three attempts), the alternate RADIUS server will be used. After 60 seconds (six attempts) local user authentication is used.



Note

This example assumes the default number of retries.

```
configure radius timeout 10
```

History

This command was first available in ExtremeXOS 10.1.

The **mgmt-access** and **netlogin** keywords were added in ExtremeXOS 11.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure radius-accounting retries

```
configure radius-accounting {mgmt-access [primary | secondary] |  
  netlogin [primary | secondary] | index} retries retries
```

Description

This command is used to set the number of retries the switch will attempt. This value may be global or on a per server basis.

Syntax Description

mgmt-access	<i>RADIUS</i> authentication for management access
netlogin	RADIUS authentication for netlogin access.
primary	Primary server.
secondary	Secondary server.
<i>index</i>	RADIUS server index.
retries	RADIUS server retries.
<i>retries</i>	RADIUS sever retries. Range 1-20.

Default

The default value is 3, with a range of 0-10.

Usage Guidelines

None.

History

This command was first available in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure radius-accounting server client-ip

```
configure radius-accounting { mgmt-access | netlogin } [ primary  
  | secondary | index ] server [ host_ipaddr | host_ipV6addr |  
  hostname] {udp_port | tls {tls_port}} client-ip [ client_ipaddr |  
  client_ipV6addr] {vr vr_name} {shared-secret {encrypted} secret}
```

Description

Configures the *RADIUS* accounting server.

Syntax Description

mgmt-access	Specifies the RADIUS authentication server for switch management.
netlogin	Specifies the RADIUS authentication server for network login.
primary	Configures the primary RADIUS authentication server.
secondary	Configures the secondary RADIUS authentication server.
<i>index</i>	RADIUS server index. Range: 1 - 2147483641.
<i>ipaddress</i>	The IP address of the server being configured.
<i>host_ipV6addr</i>	Server IPv6 address.
<i>hostname</i>	The host name of the server being configured.
<i>udp_port</i>	The UDP port to use to contact the RADIUS authentication server.
tls	Specifies using Transfer Layer Security (TLS).
<i>tls_port</i>	The TLS port to use to contact the RADIUS accounting server.
<i>ipaddress</i>	The IP address used by the switch to identify itself when communicating with the RADIUS authentication server.
<i>client_ipV6addr</i>	Client IPv6 address.
<i>vr_name</i>	Specifies the virtual router on which the client IP is located. Note: User-created VRs are supported only on the platforms listed for this feature in the Switch Engine 32.3 Feature License Requirements document.
shared-secret	Shared secret
<i>secret</i>	Secret string.
encrypted	Password is encrypted.

Default

The following lists the default behavior of this command:

- The UDP port setting is 1813.
- The TLS port setting is 2083.
- The virtual router used is *VR-Mgmt*, the management virtual router.
- Switch management and network login use the same RADIUS accounting server.

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.

Beginning with ExtremeXOS 11.2, you can specify one pair of RADIUS accounting servers for switch management and another pair for network login. To specify RADIUS accounting servers for switch management (Telnet, SSH, and console sessions), use the **mgmt-access** keyword. To specify RADIUS accounting servers for network login, use the **netlogin** keyword. If you do not specify a keyword, switch management and network login use the same pair of RADIUS accounting servers.

If you are running ExtremeXOS 11.1 or earlier and upgrade to ExtremeXOS 11.2, you do not lose your existing RADIUS accounting server configuration. Both switch management and network login use the RADIUS accounting server specified in the older configuration.

Example

The following example configures RADIUS accounting on host radius1 using the default UDP port (1813) for use by the RADIUS client on switch 10.10.20.30 using a virtual router interface of *VR-Default* for both management and network login:

```
configure radius-accounting primary server radius1 client-ip 10.10.20.30 vr vr-Default
```

The following example configures RADIUS accounting for network login on host netlog1 using the default UDP port for use by the RADIUS client on switch 10.10.20.31 using the default virtual router interface:

```
configure radius-accounting netlogin primary server netlog1 client-ip 10.10.20.31
```

History

This command was first available in ExtremeXOS 10.1.

The **mgmt-access** and **netlogin** keywords were added in ExtremeXOS 11.2.

The *index*, *host_ipV6addr*, *client_ipV6addr*, **shared-secret**, and **encrypted** keywords were added in ExtremeXOS 16.1.

The **tls** keyword with *tls_port* variable was added in ExtremeXOS 31.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure radius-accounting shared-secret

```
configure radius-accounting [primary | secondary index] shared-secret  
    {encrypted encrypted_secret | secret }
```

Description

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

Syntax Description

mgmt-access	Specifies the switch management RADIUS accounting server.
netlogin	Specifies the network login RADIUS accounting server.
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 string is already encrypted.
<i>secret</i>	The string to be used for authentication. Maximum length of 32 characters.

Default

Unconfigured.

Usage Guidelines

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

The mgmt-access keyword specifies the RADIUS accounting server used for switch management.

The netlogin keyword specifies the RADIUS accounting server used for network login.

If you do not specify the mgmt-access or netlogin keywords, the secret applies to both the primary or secondary switch management and netlogin RADIUS accounting servers.

The encrypted keyword is primarily for the output of the [show configuration](#) command, so the shared secret is not revealed in the command output. Do not use it to set the shared secret.

Example

The following command configures the shared secret as “purpleaccount” on the primary RADIUS accounting server for both management and network login:

```
configure radius primary shared-secret purpleaccount
```

The following command configures the shared secret as “greenaccount” on the primary management RADIUS accounting server:

```
configure radius mgmt-access primary shared-secret greenaccount
```

History

This command was first available in ExtremeXOS 10.1.

The encrypted keyword was added in ExtremeXOS 11.0.

The mgmt-access and netlogin keywords were added in ExtremeXOS 11.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure radius-accounting timeout

```
configure radius-accounting {mgmt-access {primary | secondary} |  
    netlogin {primary | secondary} | index } timeout sec
```

Description

Configures the timeout interval for *RADIUS*-Accounting authentication requests.

Syntax Description

mgmt-access	Specifies the switch management RADIUS authentication server.
netlogin	Specifies the network login RADIUS authentication server.
primary	Primary server.
secondary	Secondary server.
<i>index</i>	RADIUS server index.
<i>seconds</i>	Specifies the number of seconds for authentication requests. Range is 1 to 240 seconds.

Default

The default is 3 seconds.

Usage Guidelines

This command configures the timeout interval for RADIUS-Accounting authentication requests. When the timeout has expired, another authentication attempt will be made. After three failed attempts to authenticate, the alternate server will be used.

The mgmt-access keyword specifies the RADIUS accounting server used for switch management.

The netlogin keyword specifies the RADIUS accounting server used for network login.

If you do not specify the mgmt-access or netlogin keywords, the timeout interval applies to both switch management and netlogin RADIUS accounting servers.

Example

This example configures the timeout interval for RADIUS-Accounting authentication to 10 seconds. After 30 seconds (three attempts), the alternate RADIUS server will be used:



Note

This example assumes the default number of retries of 3.

```
configure radius-accounting timeout 10
```

History

This command was first available in ExtremeXOS 10.1.

The mgmt-access and netlogin keywords were added in ExtremeXOS 11.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure radius dynamic-authorization server client-ip

```
configure radius dynamic-authorization index [nas-ip [ignore | require]
| server [host_ipaddr | host_ipV6addr | hostname] {tls {tls_port}}
client-ip [client_ipaddr | client_ipV6addr] {vr vr_name} {shared-
secret {encrypted} secret}
```

Description

This command configures up to eight *RADIUS* servers with dynamic authorization.



Note

It is recommended to enable loopback mode on the *VLAN* associated with RADIUS if the RADIUS connectivity is established using a front panel port on a SummitStack.

Syntax Description

dynamic-authorization	Specifies RADIUS dynamic authorization.
<i>index</i>	RADIUS server index. Range: 1-2147483641.
nas-ip	Specifies configuring the Network Access Server (NAS) IP Address requirement.
<i>ignore</i>	Specifies to ignore the NAS-IP Address requirement.
<i>require</i>	Specifies to require the NAS-IP Address (default).
server <i>host_ipaddr</i> <i>host_ipV6addr</i>	Server IPv4 address in either IPv4 (<i>host_ipaddr</i>) or IPv6 (<i>host_ipV6addr</i>) format.
<i>hostname</i>	The host name of the server being configured.

tls	Specifies using Transmission Control Protocol (TCP).
<i>tls_port</i>	The TLS port to use to contact the RADIUS authentication server.
client-ip <i>client_ipaddr</i> <i>client_ipV6addr</i>	Client address in either IPv4 (<i>client_ipaddr</i>) or IPv6 format (<i>client_ipV6addr</i>).
vr <i>vr_name</i>	Specifies the virtual router on which the client IP is located. Note: User-created VRs are supported only on the platforms listed for this feature in the Switch Engine 32.3 Feature License Requirements document.
shared-secret	Shared secret.
<i>secret</i>	Secret string.
encrypted	Password is encrypted.

Default

The virtual router used is *VR-Mgmt*, the management virtual router.

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 example configures a RADIUS dynamic authorization server with server index 100 on host "radius1" using the default UDP port (1812) for use by the RADIUS client on switch 10.10.20.30 using a virtual router interface of *VR-Default*:

```
configure radius dynamic-authorization 100 server radius1 client-ip 10.10.20.30 vr vr-Default
```

History

This command was first available in ExtremeXOS 22.1.

The **nas-ip** option with *ignore* and *require* variables were introduced in ExtremeXOS 31.3.

The **tls** keyword with *tls_port* variable was added in ExtremeXOS 31.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure radius tls ocs

```
configure radius tls ocs [ on | off]
```

Description

This command configures Online Certificate Status Protocol (OCSP) globally for all RADIUS TLS servers.

Syntax Description

tls	Specifies Transport Layer Security (TLS).
ocs	Specifies the OCSP attribute.
on	Specifies turning ON OCSP for all RADIUS TLS servers.
off	Specifies turning OFF OCSP for all RADIUS TLS servers.

Default

ON.

Usage Guidelines

This is not configurable per server.

Example

The following example turns off OCSP for all RADIUS TLS servers:

```
# configure radius tls ocs off
```

History

This command was first available in ExtremeXOS 31.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure radius tls ocs nonce

```
configure radius tls ocs nonce [ on | off]
```

Description

Enables or disables Online Certificate Status Protocol (OCSP) nonce for RADIUS TLS servers.

Syntax Description

tls	Specifies Transport Layer Security (TLS).
ocsp	Specifies the OCSP attribute.
nonce	Specifies to cryptographically bind an OCSP request and an OCSP response with the extension <code>id-pkix-ocsp-nonce</code> to prevent replay attacks.
on	Specifies to include the <code>id-pkix-ocsp-nonce</code> extension in the OCSP request and response.
off	Specifies to exclude the extension (default).

Default

Off.

Usage Guidelines

Example

The following example configures nonce:

```
# configure radius tls ocsp nonce on
```

History

This command was first available in ExtremeXOS 32.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure radius tls ocsp override

```
configure radius tls ocsp override [url | none]
```

Description

This command configures one HTTP Online Certificate Status Protocol (OCSP) override URL for RADIUS TLS servers.

Syntax Description

tls	Specifies Transport Layer Security (TLS).
ocsp	Specifies the OCSP attribute.

override	Specifies to override the OCSP server in the AuthorityInformationAccess section of a RADIUS TLS server's certificate.
<i>url</i>	Specifies the URL of the OCSP override server. Default port is 80.
none	Specifies to remove the OCSP override URL configuration (default).

Default

None.

Usage Guidelines

Only HTTP is supported with either FQDN or IP.

Example

The following example configures an override URL of `http://radiusocsp:2021`:

```
# configure radius tls ocsp override http://radiusocsp:2021
```

History

This command was first available in ExtremeXOS 32.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure radius tls ocsp signer

```
configure radius tls ocsp signer ocsp-nocheck [on | off]
```

Description

Enables or disables Online Certificate Status Protocol (OCSP) signer's `ocsp-nocheck` for RADIUS TLS servers.

Syntax Description

tls	Specifies Transport Layer Security (TLS).
ocsp	Specifies the OCSP attribute.
ocsp-nocheck	Specifies the extension <code>id-pkix-ocsp-nocheck</code> . If present in the OCSP signer's certificate, then it is trusted for its lifetime.
on	Specifies to override the <code>id-pkix-ocsp-nocheck</code> extension in the OCSP signer's certificate and forces the extension as if it is present.

off	Specifies to behave per the extension's presence in the OCSP signer's certificate. If not present and the OCSP signer is not root CA, then the whole OCSP will fail (default).
signer	Specifies the OCSP signer that signs the OCSP response.

Default

Off.

Usage Guidelines

Example

The following example enables OCSP signer's nocheck for a RADIUS TLS server.

```
# configure radius tls ocsp signer ocsp-nocheck on
```

History

This command was first available in ExtremeXOS 32.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure radius tls tcp-user-timeout

```
configure radius tls tcp-user-timeout [ seconds | default]
```

Description

This command configures TCP connection timeout globally for all TLS servers.

Syntax Description

tls	Specifies Transport Layer Security (TLS).
tcp-user-timeout	Specifies the TCP user timeout attribute.
<i>seconds</i>	Specifies the timeout in seconds.
default	Specifies to use the system default timeout.

Default

Use the system's TCP user timeout setting.

Usage Guidelines

This is not configurable per server.

Example

The following example sets the TCP user timeout to 60 seconds.

```
configure radius tls tcp-user-timeout 60
```

History

This command was first available in ExtremeXOS 31.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure rip add vlan

```
configure rip add vlan [vlan_name | all]
```

Description

Configures *RIP* on an IP interface.

Syntax Description

<i>vlan_name</i>	Specifies a <i>VLAN</i> name.
all	Specifies all VLANs.

Default

N/A.

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:

```
# configure rip add finance
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms with a Base license or higher as described in the [*Switch Engine 32.3 Feature License Requirements*](#) document.

configure rip delete vlan

```
configure rip delete vlan [vlan_name | all]
```

Description

Disables *RIP* on an IP interface.

Syntax Description

<i>vlan_name</i>	Specifies a <i>VLAN</i> name.
all	Specifies all VLANs.

Default

N/A.

Usage Guidelines

When an IP interface is created, RIP configuration is disabled on the interface by default. When the RIP interface is disabled by this command, the parameters are not reset to default automatically.

Example

The following command deletes RIP on a VLAN named finance:

```
# configure rip delete finance
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms with a Base license or higher as described in the [*Switch Engine 32.3 Feature License Requirements*](#) document.

configure rip garbagetime

```
configure rip garbagetime {seconds}
```

Description

Configures the *RIP* garbage time.

Syntax Description

<i>seconds</i>	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:

```
# configure rip garbagetime 60
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms with a Base license or higher as described in the [***Switch Engine 32.3 Feature License Requirements***](#) document.

configure rip import-policy

```
configure rip import-policy [policy-name | none]
```

Description

Configures the import policy for *RIP*.

Syntax Description

<i>policy-name</i>	Specifies the policy.
--------------------	-----------------------

Default

No policy.

Usage Guidelines

An import policy is used to modify route attributes while adding RIP routes to the IP route table. The import policy cannot be used to determine the routes to be added to the routing table.

Use the none option to remove an import policy.

Example

The following example applies the policy campuseast to RIP routes:

```
# configure rip import-policy campuseast
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms with a Base license or higher as described in the [***Switch Engine 32.3 Feature License Requirements***](#) document.

configure rip routetimeout

```
configure rip routetimeout seconds
```

Description

Configures the route timeout period.

Syntax Description

<i>seconds</i>	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:

```
# configure rip routetimeout 120
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms with a Base license or higher as described in the [***Switch Engine 32.3 Feature License Requirements***](#) document.

configure rip updatetime

```
configure rip updatetime seconds
```

Description

Specifies the time interval in seconds within which *RIP* sends update packets.

Syntax Description

<i>seconds</i>	Specifies a time in seconds. The range is 10 to 180.
----------------	--

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. Timer minimum is 10 seconds and maximum is 180 seconds.

Example

The following command sets the update timer to 60 seconds:

```
# configure rip updatetime 60
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms with a Base license or higher as described in the [*Switch Engine 32.3 Feature License Requirements*](#) document.

configure rip vlan cost

```
configure rip vlan [vlan_name | all] cost cost
```

Description

Configures the cost (metric) of the interface.

Syntax Description

<i>vlan_name</i>	Specifies a <i>VLAN</i> name.
all	Specifies all VLANs.
<i>cost</i>	Specifies a cost metric.

Default

The default setting is 1.

Usage Guidelines

The specified interface cost is added to the cost of the route received through this interface.

Example

The following command configures the cost for the VLAN finance to a metric of 3:

```
# configure rip vlan finance cost 3
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms with a Base license or higher as described in the [*Switch Engine 32.3 Feature License Requirements*](#) document.

configure rip vlan route-policy

```
configure rip vlan [vlan_name | all] route-policy [in | out] [policy-  
name | none]
```

Description

Configures *RIP* to ignore certain routes received from its neighbor, or to suppress certain routes when performing route advertisements.

Syntax Description

<i>vlan_name</i>	Specifies a <i>VLAN</i> name.
all	Specifies all VLANs.
<i>policy-name</i>	Specifies a policy.
none	Removes any policy from the VLAN.

Default

N/A.

Usage Guidelines

Use the in option to configure an input route policy, which determines 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.

Use the out option to configure an output route policy, which determines which RIP routes are advertised on the VLAN.

Example

The following command configures the VLAN backbone to accept selected routes from the policy nosales:

```
# configure rip vlan backbone route-policy in nosales
```

The following command uses the policy nosales to determine which RIP routes are advertised into the VLAN backbone:

```
# configure rip vlan backbone route-policy out nosales
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms with a Base license or higher as described in the [*Switch Engine 32.3 Feature License Requirements*](#) document.

configure rip vlan rxmode

```
configure rip [vlan vlan_name | all] rxmode [none | v1only | v2only |
any]
```

Description

Syntax Description

<i>vlan_name</i>	Specifies to apply settings to specific <i>VLAN</i> name.
all	Specifies all VLANs.
none	Specifies to drop all received <i>RIP</i> 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.

Default

N/A.

Usage Guidelines

None.

Example

The following command configures the receive mode for the VLAN finance to accept only RIP version 1 format packets:

```
# configure rip finance rxmode v1only
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms with a Base license or higher as described in the [*Switch Engine 32.3 Feature License Requirements*](#) document.

configure rip vlan trusted-gateway

```
configure rip vlan [vlan_name | all] trusted-gateway [policy-name |
none]
```

Description

Configures a trusted neighbor policy to determine trusted *RIP* router neighbors for the *VLAN* on the switch running RIP.

Syntax Description

<i>vlan_name</i>	Specifies a VLAN name.
all	Specifies all VLANs.
<i>policy-name</i>	Specifies a policy.
none	Removes any trusted-gateway policy from the VLAN.

Default

N/A.

Usage Guidelines

Use this command to set a policy to determine trusted neighbors. A neighbor is defined by its IP address. Only the RIP control packets from trusted neighbors will be processed.

Example

The following command configures RIP to use the policy nointernet to determine from which RIP neighbor to receive (or reject) the routes to the VLAN backbone:

```
# configure rip vlan backbone trusted-gateway nointernet
```

History

This command is available on all platforms with a Base license or higher as described in the *Switch Engine 32.3 Feature License Requirements* document.

Platform Availability

This command is available on all platforms with an Edge, Advanced Edge, or Core license.

configure rip vlan txmode

```
configure rip [vlan vlan_name | all] txmode [none | v1only | v1comp |  
          v2only]
```

Description

Changes the *RIP* transmission mode for one or all *VLANs*.

Syntax Description

<i>vlan_name</i>	Specifies to apply settings to a specific VLAN name.
all	Specifies all VLANs.
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.

Default

N/A.

Usage Guidelines

None.

Example

The following command configures the transmit mode for the VLAN finance to transmit version 2 format packets to the broadcast address:

```
# configure rip finance txmode v1comp
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms with a Base license or higher as described in the [*Switch Engine 32.3 Feature License Requirements*](#) document.

configure ripng add

```
configure ripng add [vlan vlan-name | tunnel tunnel-name | [vlan |  
                  tunnel] all]
```

Description

Configures *RIPng* on an IP interface.

Syntax Description

<i>vlan-name</i>	Specifies an IPv6 configured <u>VLAN</u> .
<i>tunnel-name</i>	Specifies an IPv6 tunnel.
all	Specifies all IPv6 configured VLANs or tunnels.

Default

N/A.

Usage Guidelines

For RIPng to be active on the interface, it must also be globally enabled using the command `disable ripng export [direct | ospfv3 | ospfv3-extern1 | ospfv3-extern2 | ospfv3-inter | ospfv3-intra | static | isis | isis-level-1 | isis-level-1-external | isis-level-2 | isis-level-2-external | bgp]`. If the keyword **all** is specified, all IPv6 configured VLANs or tunnels will be configured for RIPng.

Example

The following command configures RIPng on the VLAN finance:

```
configure ripng add finance
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Edge, Advanced Edge, or Core license. For licensing information, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure ripng cost

```
configure ripng [vlan vlan-name | tunnel tunnel-name] cost metric
```

Description

Configures the cost (metric) of the interface..

Syntax Description

<i>vlan-name</i>	Specifies an IPv6 configured <u>VLAN</u> .
<i>tunnel-name</i>	Specifies an IPv6 tunnel.
<i>metric</i>	Specifies a cost metric. Range is 1 to 15.

Default

The default setting is 1.

Usage Guidelines

The specified interface cost is added to the cost of the route received through this interface.

Example

The following command configures the cost for the VLAN finance to a metric of 3:

```
configure ripng vlan finance cost 3
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Edge, Advanced Edge, or Core license. For licensing information, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure ripng delete

```
configure ripng delete [vlan vlan-name | tunnel tunnel-name | [vlan | tunnel] all]
```

Default

Removes an interface from RIPng routing.

Syntax Description

<i>vlan-name</i>	Specifies an IPv6 configured <u>VLAN</u> .
<i>tunnel-name</i>	Specifies an IPv6 tunnel.
all	Specifies all IPv6 configured VLANs or tunnels.

Default

N/A.

Usage Guidelines

This command removes an interface from RIPng routing. However, the RIPng-specific interface configuration will be preserved, even if RIPng is unconfigured on the interface. The interface

configuration information is removed only when the IPv6 interface itself gets deleted by, for example, by unconfiguring all the IPv6 addresses on the interface.

Example

The following command removes the VLAN finance from RIPng routing:

```
configure ripng delete finance
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Edge, Advanced Edge, or Core license. For licensing information, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure ripng garbagetime

```
configure ripng garbagetime {seconds}
```

Description

Configures the *RIPng* garbage time.

Syntax Description

<i>seconds</i>	Specifies a time in seconds. Range is 10 to 2400 seconds.
----------------	---

Default

120 seconds.

Usage Guidelines

This command configures the time interval after which a route in the RIPng routing database that has expired will be removed. The value is rounded off to nearest multiple of 10.

Example

The following command configures the RIPng garbage time to have a 60-second delay:

```
configure ripng garbagetime 60
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Edge, Advanced Edge, or Core license. For licensing information, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure ripng import-policy

```
configure ripng import-policy [policy-name | none]
```

Description

Configures the import policy for *RIPng*.

Syntax Description

<i>policy-name</i>	Specifies the policy.
--------------------	-----------------------

Default

No policy.

Usage Guidelines

Use this command to configure the policy to be applied to RIPng routes installed into the system routing table from the RIPng routing process. This policy can be used to modify parameters associated with routes installed into the routing table. The import policy cannot be used to determine the routes to be added to the routing table.

Use the none option to remove the import policy.

The following is a sample policy file that can be used with RIPng. It changes the metric to 12 for any routes from the subnets 2001:db8:2ccc::/64 and 2001:db8:2ccd::/64:

```
entry filter_routes {
    If match any{
        nlri 2001:db8:2ccc:: /64;
        nlri 2001:db8:2ccd:: /64;
    }
    then {
        cost 12;
    }
}
```

Example

The following example applies the policy `campuseast` to RIPng routes:

```
configure ripng import-policy campuseast
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Edge, Advanced Edge, or Core license. For licensing information, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure ripng route-policy

```
configure ripng [vlan vlan-name | tunnel tunnel-name] route-policy [in |  
out] [policy-name | none]
```

Description

Configures [RIPng](#) to ignore or modify certain routes received from its neighbors, or to suppress certain routes when performing route advertisements.

Syntax Description

<i>vlan-name</i>	Specifies an IPv6 configured VLAN .
<i>tunnel-name</i>	Specifies an IPv6 tunnel.
<i>policy-name</i>	Specifies a policy.
none	Removes any policy from the VLAN.

Default

N/A.

Usage Guidelines

Use the `in` option to configure an input route policy, which determines which RIPng routes are accepted as valid routes from RIPng neighbors. This policy can be combined with the trusted neighbor policy to accept selected routes only from a set of trusted neighbors.

Use the `out` option to configure an output route policy, which determines which RIPng routes are advertised to other RIPng neighbors.

The following is a sample policy file that could be used with RIPng. It will drop any routes from the subnets 2001:db8:2ccc::/64 and 2001:db8:2ccd::/64:

```
entry filter_routes {
  If match any{
    nlri 2001:db8:2ccc:: /64;
    nlri 2001:db8:2ccd:: /64;
  }
  then {
    deny;
  }
}
```

Example

The following command configures the VLAN backbone to accept routes from its neighbor as specified by the policy nosales:

```
configure ripng vlan backbone route-policy in nosales
```

The following command uses the policy nosales to determine which *RIP* routes are advertised into the VLAN backbone:

```
configure rip vlan backbone route-policy out nosales
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Edge, Advanced Edge, or Core license. For licensing information, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure ripng routetimeout

```
configure ripng routetimeout seconds
```

Description

Configures the route timeout period for *RIPng*.

Syntax Description

<i>seconds</i>	Specifies a time in seconds. Range is 10 to 3600.
----------------	---

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.

The configured value is rounded off to the nearest multiple of 10.

Example

The following example sets the route timeout period to 120 seconds:

```
configure ripng routetimeout 120
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Edge, Advanced Edge, or Core license. For licensing information, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure ripng trusted-gateway

```
configure ripng [vlan vlan-name | tunnel tunnel-name] trusted-gateway
                [policy-name | none]
```

Description

Configures a trusted neighbor policy to determine trusted [RIPng](#) router neighbors for the interfaces on the switch running RIPng.

Syntax Description

<i>vlan-name</i>	Specifies an IPv6 configured VLAN .
<i>tunnel-name</i>	Specifies an IPv6 tunnel.
<i>policy-name</i>	Specifies a policy.
none	Removes any trusted-gateway policy from the VLAN.

Default

None. Control packets from all of the neighbors are processed.

Usage Guidelines

Use this command to set a policy to determine trusted neighbors. A neighbor is defined by its IP address. Only the RIPng control packets from trusted neighbors will be processed.

The following policy designates neighbors from the fe80:202:b3ff:fe4a:6ada:: /64 subnet and the neighbor at fe80:203::b3ff:fe4a:6ada as trusted gateways:

```
entry filter_gateways {
  If match any{
    nlri fe80:202:b3ff:fe4a:6ada:: /64;
    nlri fe80:203::b3ff:fe4a:6ada:: /64;
  }
  then {
    permit;
  }
}
```

Example

The following command configures RIPng to use the policy nointernet to determine from which RIPng neighbor to receive (or reject) the routes to the VLAN backbone:

```
configure ripng vlan backbone trusted-gateway nointernet
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Edge, Advanced Edge, or Core license. For licensing information, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure ripng updatetime

```
configure ripng updatetime seconds
```

Description

Specifies the time interval in seconds within which *RIPng* sends update packets.

Syntax Description

<i>seconds</i>	Specifies a time in seconds. The range is 10 to 3600.
----------------	---

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. Timer minimum is 10 second and maximum is 3600 seconds.

Example

The following command sets the update timer to 60 seconds:

```
configure ripng updatetime 60
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Edge, Advanced Edge, or Core license. For licensing information, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure switch safe-default-script

```
configure switch safe-default-script
```

Description

Allows you to change management access to your device and to enhance security.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command runs an interactive script that prompts you to choose to enable or disable SNMP, Telnet, and enabled ports.

Refer to “Using Safe Defaults Mode” in the [Switch Engine 32.3 User Guide](#) for complete information on the safe default mode.

After you issue this command, the system presents you with the following interactive script:

```
Telnet is enabled by default. Telnet is unencrypted and has been the target of security
exploits in the past. Would you like to disable Telnet? [y/N]: SNMP access is enabled
by default. SNMP uses no encryption, SNMPv3 can be configured to eliminate this problem.
Would you like to disable SNMP? [y/N]: All ports are enabled by default. In some secure
applications, it maybe more desirable for the ports to be turned off. Would you like
unconfigured ports to be turned off by default? [y/N]: Changing the default failsafe
account username and password is highly recommended. If you choose to do so, please
remember the username and password as this information cannot be recovered by Extreme
Networks. Would you like to change the failsafe account username and password now? [y/N]:
Would you like to permit failsafe account access via the management port? [y/N]: Since you
have chosen less secure management methods, please remember to increase the security of
your network by taking the following actions: * change your admin password * change your
```

```
failsafe account username and password * change your SNMP public and private strings *  
consider using SNMPv3 to secure network management traffic
```

Example

The following command reruns the interactive script to configure management access:

```
configure switch safe-default-script
```

History

This command was first available in ExtremeXOS 11.2.

The **switch** keyword was added in ExtremeXOS 30.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure security fips-mode

```
configure security fips-mode [on | off]
```

Description

This command enables you to toggle between the default OpenSSL library (FIPS compatible) and FIPS capable library.

Syntax Description

on	Enables FIPS mode.
off	Disable FIPS mode.

Default

Off.

Usage Guidelines

After enabling/disabling FIPS, EPM will be notified to change the bit dedicated to FIPS Mode. As per requirement, currently SSH and SNMP will use this bit to toggle between normal and FIPS mode.

Example

```
# sh security fips-mode  
FIPS Mode (current)      : Off  
FIPS Mode (configured)  : Off  
  
# configure security fips-mode on
```

```
FIPS mode will be enabled only after rebooting the switch.
SNMPv3 users configured with either md5 authentication or DES encryption will be
discarded after reboot.
SSH existing configuration of ciphers/MACs will be lost after reboot.
Python scripting configuration is ignored when FIPS mode is 'on'.

# show security fips-mode
FIPS Mode (current)      : On
FIPS Mode (configured)  : On
```

History

This command was first available in ExtremeXOS 21.1.

Current and configured information added in ExtremeXOS 32.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure security python

```
configure security python [on | off]
```

Description

Turns on or off external Python scripting support when FIPS mode is turned off.

Syntax Description

on	Turns on external Python scripting support (default).
off	Turns off external Python scripting support.

Default

By default, when FIPS mode is off, external Python scripting support is enabled.

Usage Guidelines

To enable external Python scripting support with the command, FIPS mode must be turned off (`configure security fips-mode [on | off]`). Python scripting configuration is ignored when FIPS mode is turned on.

Example

The following example turns off external Python scripting support:

```
# configure security python off
```

History

This command was first available in ExtremeXOS 32.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure sflow agent ipaddress

```
configure sflow agent {ipaddress} ipaddress
```

Description

Configures the sFlow agent's IP address.

Syntax Description

<i>ipaddress</i>	Specifies the IP address from which sFlow data is sent on the switch.
------------------	---

Default

The default configured IP address is 0.0.0.0, but the effective IP address is the management port IP address.

Usage Guidelines

This command allows you to configure the IP address of the sFlow agent. Typically, you would set this to the IP address used to identify the switch in the network management tools that you use. The agent address is stored in the payload of the sFlow data, and is used by the sFlow collector to identify each agent uniquely. The default configured value is 0.0.0.0, but the switch will use the management port IP address if it exists.

Both the commands `unconfigure ports monitor vlan` and `unconfigure sflow agent` will reset the agent parameter to the default.

Example

The following command sets the sFlow agent's IP address to 10.2.0.1:

```
configure sflow agent ipaddress 10.2.0.1
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure sflow collector ipaddress

```
configure sflow collector {ipaddress} ipaddress {port udp-port-number}
{vr vr_name}
```

Description

Configures the sFlow collector IP address.

Syntax Description

<i>ipaddress</i>	Specifies the IP address to send the sFlow data.
<i>udp-port-number</i>	Specifies the UDP port to send the sFlow data.
<i>vr_name</i>	Specifies from which virtual router to send the sFlow data. Note: User-created VRs are supported only on the platforms listed for this feature in the Switch Engine 32.3 Feature License Requirements document.

Default

The following values are the defaults for this command:

- UDP port number—6343
- Virtual router—[*VR-Mgmt*](#) (previously called VR-0)

Usage Guidelines

This command allows you to configure where to send the sFlow data. You must specify an IP address for the sFlow data collector, and you may specify a particular UDP port, if your collector uses a non-standard port. You may also need to specify from which virtual router to send the data.

You can configure up to four sFlow collectors. Each unique IP address/UDP port/virtual router combination identifies a collector.

Both the commands [unconfigure ports monitor vlan](#) and [unconfigure sflow collector](#) will reset the collector parameters to the default.

Example

The following command specifies that sFlow data should be sent to port 6343 at IP address 192.168.57.1 using the virtual router VR-Mgmt:

```
configure sflow collector ipaddress 192.168.57.1
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure sflow max-cpu-sample-limit

```
configure sflow max-cpu-sample-limit rate
```

Description

Configures the maximum number of sFlow samples handled by the CPU per second.

Syntax Description

<i>rate</i>	Specifies the maximum sFlow samples per second.
-------------	---

Default

The default value is 2000 samples per second.

Usage Guidelines

This command configures the maximum number of samples sent to the CPU per second. If this rate is exceeded, the internal sFlow CPU throttling mechanism kicks in to limit the load on the CPU.

Every time the limit is reached, the sample rate is halved (the value of number in the `configure sflow sample-rate number` or `configure sflow ports port_list sample-rate number` command is doubled) on the slot (SummitStack) or ports (stand-alone switch) on which maximum number of packets were received during the last snapshot.

This effectively halves the sampling frequency of all the ports on that slot or stand-alone switch with a sub-sampling factor of 1. The sampling frequency of ports on that slot or stand-alone switch with a sub-sampling factor greater than 1 will not change; the sub-sampling factor is also halved so the that the same rate of samples are sent from that port.

The maximum CPU sample rate is based on the total number of samples received from all the sources. The valid range is 100 to 200000 samples per second.

Example

The following command specifies that the sFlow maximum CPU sample rate should be set to 4000 samples per second:

```
configure sflow max-cpu-sample-limit 4000
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure sflow poll-interval

```
configure sflow poll-interval seconds
```

Description

Configures the sFlow counter polling interval.

Syntax Description

<i>seconds</i>	Specifies the number of seconds between polling each counter. The value can range from 0 to 3600 seconds.
----------------	---

Default

The default polling interval is 20 seconds.

Usage Guidelines

Each sFlow statistics counter is polled at regular intervals, and this data is then sent to the sFlow collector. This command is used to set the polling interval. To manage CPU load, polling for sFlow enabled ports are distributed over the polling interval, so that all ports are not polled at the same instant. For example, if the polling interval is 20 seconds and there are twenty counters, data is collected successively every second.

Specifying a poll interval of 0 (zero) seconds disables polling.

Example

The following command sets the polling interval to 60 seconds:

```
configure sflow poll-interval 60
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure sflow ports sample-rate

```
configure sflow ports port_list sample-rate number
```

Description

Configures the sFlow per-port sampling rate.

Syntax Description

<i>port_list</i>	Specifies a list of ports.
<i>number</i>	Specifies the fraction (1/number) of packets to be sampled.

Default

The default number is 8192, unless modified by the `configure sflow sample-rate` command.

Usage Guidelines

This command configures the sampling rate on a particular set of ports, and overrides the system-wide value set in the `configure sflow sample-rate` command. The rate is rounded off to the next power of two, so if 400 is specified, the sample rate is configured as 512. The valid range is 256 to 536870912.

All ports on the switch are sampled individually.

Example

The following command sets the sample rate for the ports 4:6 to 4:10 to one packet out of every 16384:

```
configure sflow ports 4:6-4:10 sample-rate 16384
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure sflow sample-rate

```
configure sflow sample-rate number
```

Description

Configures the sFlow default sampling rate.

Syntax Description

<i>number</i>	Specifies the fraction (1/number) of packets to be sampled.
---------------	---

Default

The default number is 8192.

Usage Guidelines

This command configures the default sampling rate. This is the rate that newly enabled sFlow ports will have their sample rate set to. Changing this rate will not affect currently enabled sFlow ports. The rate is rounded off to the next power of two, so if 400 is specified, the sample rate is configured as 512. The valid range is 256 to 536870912.

Configuring a lower number for the sample rate means that more samples will be taken, increasing the load on the switch. Do not configure the sample rate to a number lower than the default unless you are sure that the traffic rate on the source is low.

The minimum rate that these platforms sample is 1 out of every 256 packets. If you configure a rate to be less than 256, the switch automatically rounds up the sample rate to 256.

Example

The following example sets the sample rate to one packet out of every 16384:

```
configure sflow sample-rate 16384
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure sharing add ports

```
configure sharing port slot slot distribution-list [port_list | add  
port_list | all]
```

Description

Adds ports to a load-sharing, or link aggregation, group. By using link aggregation, you use multiple ports as a single logical port. Link aggregation also provides redundancy because traffic is redistributed to the remaining ports in the LAG if one port in the group goes down.

Syntax Description

<i>port</i>	Specifies the logical port for a load-sharing group or link aggregation group (LAG). This number also functions as the LAG Group ID.
<i>port_list</i>	Specifies one or more ports or slots and ports to be grouped in the LAG.
add	Adds a port list to the existing distribution port list for the given slot.
all	All active members of the group are eligible for distribution for packets received on the given slot. This is the existing behavior and the default. This option effectively deletes any existing configured port list for the slot.

Default

N/A.

Usage Guidelines

Use this command to dynamically add ports to a load-sharing group, or link aggregation group (LAG).



Note

You must create a LAG (or load-sharing group) before you can configure the LAG. To create a LAG, see [enable sharing grouping](#) on page 2302.

VMAN ports can belong to LAGs. If any port in the LAG is enabled for VMAN, all ports in the group are automatically enabled to handle jumbo size frames. Also, VMAN is automatically enabled on all ports of the untagged LAG.

To verify your configuration, use the [show ports sharing](#) command.



Note

All ports that are designated for the LAG must be removed from all VLANs prior to configuring the LAG.

ExtremeSwitching Series Switches

The following guidelines apply to link aggregation on the ExtremeSwitching series switches:

- One static LAG can contain up to 8 ports.
- An LACP LAG can include a maximum of 16 ports; out of these up to 8 can be selected links and the remaining 8 will be standby links.
- A Health Check LAG can contain up to 8 ports.

SummitStack only

The following guidelines apply to link aggregation:

- A static LAG can include a maximum of 8 ports.
- An LACP LAG can include a maximum of 16 ports; out of these up to 8 can be selected links and the remaining 8 will be standby links.

- A Health Check LAG can include a maximum of 8 ports.

Example

The following example adds port 3 to the LAG with the logical port 4 on a switch:

```
configure sharing 3 add port 4
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure sharing address-based custom

```
configure sharing address-based custom [ipv4 [source-only | destination-only | source-and-destination] | hash-algorithm [xor | crc-16 | crc-32 [lower | upper]]]
```

Description

This command configures the part of the packet examined by the switch when selecting the egress port for transmitting link aggregation, or load-sharing, data.

Syntax Description

ipv4	IPv4 hash configuration for custom load sharing and L2VPN sharing.
source-only	Indicates that the switch should examine the IP source address only.
destination-only	Indicates that the switch should examine the IP destination address only.
source-and-destination	Indicates that the switch should examine the IP source and destination address.
hash-algorithm	Hash algorithm for custom load sharing and L2VPN sharing.
xor	Use exclusive-OR for load sharing hash computation.
crc-16	Use CRC-16 for load sharing hash computation.
crc-32	Use CRC-32 for load sharing hash computation.
lower	Use lower 16 bits of CRC32 for load sharing hash computation.
upper	Use upper 16 bits of CRC32 for load sharing hash computation.

Default

Algorithm: source-and-destination.

Hash algorithm: xor.

Usage Guidelines

This command specifies the part of the packet header that the switch examines to select the egress port for address-based load-sharing trunks. The address-based load-sharing setting is global and applies to all load-sharing trunks, or LAGs, that are address-based and configured with a custom algorithm. You change this setting by issuing the command again with a different option.

The addressing information examined is based on the packet protocol as follows:

- IPv4 packets—Uses the source and destination IPv4 addresses and Layer4 port numbers as specified with this command.
- IPv6 packets—Uses the source and destination IPv6 addresses and Layer4 port numbers.
- MPLS packets—Uses the top, second, and reserved labels and the source and destination IP addresses.
- Non-IP Layer 2—Uses the VLAN ID, the source and destination MAC addresses, and the ethertype.

The xor hash algorithm guarantees that the same egress port is selected for traffic distribution based on a pair of IP addresses, Layer4 ports, or both, regardless of which is the source and which is the destination.

For IP-in-IP and GRE tunneled packets, the switch examines the inner header to determine the egress port.

To verify your configuration, use the `show ports sharing` command.

Example

The following example configures the switch to examine the source IP address:

```
# configure sharing address-based custom ipv4 source-only
```

History

This command was first available in ExtremeXOS 12.3.

Platform Availability

This command is supported on the ExtremeSwitching 5520 platform.

configure sharing address-based custom hash-seed

```
configure sharing address-based custom hash-seed [seed | switch-mac-  
address]
```

Description

Configures the hash seed used in the CRC hashing algorithms of the “custom” load sharing algorithm.

Syntax Description

address-based	Selects address-based sharing.
custom	Configuration for address-based custom load sharing and L2VPN sharing.
hash-seed	Selects configuring hash seed used with CRC hash algorithms.
<i>seed</i>	Sets the hash seed value. Prior to ExtremeXOS22.5, the default value was 0x7F2193EA.
switch-mac-address	Use the last four bytes of the switch's MAC address to create a unique seed value (Default).

Default

The default is **switch-mac-address**.

Usage Guidelines

The default configuration of the hash seed is **switch-mac-address**, which uses the last four bytes of the switch's MAC address as the hash seed to provide a unique seed value on all Extreme Networks switches in the network. Such a configuration prevents hash polarization in MLAG network configurations by default.

Prior to supporting configuring the hash seed (ExtremeXOS 30.1), the default value of the hash seed was 0x7F2193EA. You can restore the legacy default behavior for the hash seed by explicitly configuring this legacy value.

To verify your hash seed configuration, use the `show ports port_list sharing distribution configuration` or `show {port port_number} sharing {detail}` commands.

Example

The following example sets the hash seed value to "123456789":

```
configure sharing address-based custom hash-seed 123456789
```

History

This command was first available in ExtremeXOS 30.1.

Platform Availability

This command is supported on the ExtremeSwitching 5520 platform.

configure sharing algorithm

```
configure sharing master_port algorithm [address-based [L2 | L3 | L3_L4
| custom] | port-based]
```

Description

Modifies the distribution algorithm of an existing [LAG](#).

Syntax Description

<i>master_port</i>	Specifies the master logical port for the load-sharing group or LAG.
algorithm	Specifies modifying the distribution algorithm of an existing LAG.
address-based	Specifies link aggregation by address-based algorithm.
L2	Specifies address-based link aggregation by Layer 2. This is the default.
L3	Specifies address-based link aggregation by Layer 3.
L3_L4	Specifies address-based link aggregation by Layer 3 IP plus Layer 4 port.
custom	Selects the custom link aggregation algorithm configured with the following command: <code>configure sharing address-based custom [ipv4 [L3-and-L4 source-only destination-only source-and-destination] hash-algorithm [xor crc-16]]</code> . The configuration of the custom option applies to all LAGs on the switch.
port-based	Selects port-based load sharing groups.

Default

Address-based link aggregation by Layer 2 is the default.

Usage Guidelines

This command allows you to modify the distribution algorithm of an existing LAG, created using the command [enable sharing grouping](#) on page 2302.

If you select the **custom** option, you configure the customer link aggregation algorithm with the following command: `configure sharing address-based custom [ipv4 [L3-and-L4 | source-only | destination-only | source-and-destination] | hash-algorithm [xor | crc-16]]`

Since the custom and port-based algorithms may not be used at the same time, changing the algorithm on multiple groups between the custom and port-based algorithms requires changing the algorithm on these groups to either L2, L3, or L3_L4 as an intermediate step.

Example

The following example sets the distribution algorithm for the LAG on port 24 to address-based link aggregation by Layer 3 IP plus Layer 4 port:

```
# configure sharing 24 algorithm address-based L3_L4
```

History

This command was first available in ExtremeXOS 22.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure sharing delete ports

```
configure sharing port slot slot distribution-list [port_list | delete
port_list | all]
```

Description

Deletes ports from a link aggregation, or load-sharing, group.

Syntax Description

<i>port</i>	Specifies the logical port for a load-sharing group or a LAG . This number also functions as the LAG Group ID.
<i>port_list</i>	Specifies one or more ports or slots and ports to be grouped in the LAG.

Default

N/A.

Usage Guidelines

Use this command to dynamically delete ports from a load-sharing group, or link aggregation group (LAG). This command applies to static and dynamic link aggregation.

Example

The following example deletes port 3:12 from the LAG with the logical port, or LAG Group ID, 3:9:

```
configure sharing 3:9 delete port 3:12
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure sharing distribution-mode

```
configure sharing master_port distribution-mode [all | local-slot | port-lists]
```

Description

This command provides two different configuration options for specifying subsets of active member ports as eligible for distribution. Both of these options specify a subset of the active member ports on a per slot basis. The specific choice of configuration is described in the CLI as a “distribution-mode”. The choice of distribution mode is configurable per [LAG](#).

Syntax Description

all	All active members of the group are eligible for distribution on all slots in the switch. This is the existing behavior and the default.
local-slot	If there are one or more active members of the group on the slot where traffic is received, distribution will be restricted to these “local-slot” members.
port-lists	If there are one or more active members of the group in the configured distribution port list for the slot on which traffic is received, distribution will be restricted to these configured ports.

Default

All.

Usage Guidelines

The “local-slot” distribution mode restricts distribution of unicast packets to the active LAG members on the same slot where the packet was received. If no active LAG members are present on the slot where the packet was received, all active LAG member ports are included in the distribution algorithm. The “local-slot” distribution mode may be specified during LAG creation with the “enable sharing” CLI command. It may also be configured dynamically with the “configure sharing” command. This distribution mode is self-configuring in the sense that no configuration is required other than the specification of the “local-slot” distribution mode. Addition or deletion of LAG member ports via the “configure sharing <master_port> [add | delete] <port_list>” command is automatically handled. The “local-slot” distribution mode is useful for reducing the fabric bandwidth load of a switch.

Example

```
# show sharing distribution configuration
Config  Distribution Distribution
Master  Mode       Lists
=====
1:1      Port Lists   Slot 1: 1:1-10, 1:15
                               Slot 5: 1:11-22
1:25     Local Slot  Slot 1: 1:25
                               Slot 5: 1:26
5:1      Port Lists
```



```
5:10      All      Slot 1: 5:11
           Slot 5: 5:10
```

History

This command was first available in ExtremeXOS 21.1.

Platform Availability

This command is available on SummitStack switches.

configure sharing health-check member-port add tcp-tracking

```
configure sharing health-check member-port port add tcp-tracking IP
      Address {tcp-port TC Port frequency sec misses count}
```

Description

Configures monitoring for each member port of a health check [LAG](#).

Syntax Description

<i>port</i>	Specifies the member port.
<i>IP Address</i>	Specifies the IP address to monitor.
<i>TCP Port</i>	Specifies the TCP port to watch. The default is port 80.
<i>sec</i>	Specifies the frequency in seconds at which tracking takes place. The default is 10 seconds.
<i>count</i>	Specifies the number of misses before a connection loss is reported. The default is 3 misses.

Default

N/A.

Usage Guidelines

To configure a health check LAG, you first create a health check type of LAG using the [enable sharing grouping](#) command. Then use this command to configure the monitoring for each member port. You can configure each member port to track a particular IP address, but only one IP address per member port.

To display the monitoring configuration for a health check LAG, use the [show sharing health-check](#) command.

To display the link aggregation configured on a switch, use the [show ports sharing](#) command.

Example

The following commands configure four different member ports:

```
# configure sharing health-check member-port 10 add track-tcp 10.1.1.1 tcp-port 23
# configure sharing health-check member-port 11 add track-tcp 10.1.1.2 tcp-port 23
# configure sharing health-check member-port 12 add track-tcp 10.1.1.3
# configure sharing health-check member-port 13 add track-tcp 10.1.1.4
```

When the TCP port, seconds, or counts are not specified, they default to the values described in the Syntax Description.

History

This command was first available in ExtremeXOS 12.1.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure sharing health-check member-port delete tcp-tracking

```
configure sharing health-check member-port port delete tcp-tracking IP
Address {tcp-port TC Port}
```

Description

Unconfigures monitoring for each member port of a health check LAG.

Syntax Description

<i>port</i>	Specifies the member port.
<i>IP Address</i>	Specifies the IP address.
<i>TCP Port</i>	Specifies the TCP port.

Default

N/A.

Usage Guidelines

Use this command to remove the monitoring configuration on the ports of a health check link aggregation group. Each port must be unconfigured separately, specifying the IP address and TCP port.

Example

The following command removes the configuration setting on port 12 that monitors IP address 10.1.1.3:

```
# configure sharing health-check member-port 12 delete track-tcp 10.1.1.3
```

History

This command was first available in ExtremeXOS 12.1.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure sharing health-check member-port tcp-tracking

```
configure sharing health-check member-port port [disable | enable] tcp-  
tracking
```

Description

Enables or disables configured monitoring on a member port of a health check LAG.

Syntax Description

<i>port</i>	Specifies the member port.
-------------	----------------------------

Default

N/A.

Usage Guidelines

This disables/enables monitoring on a particular member port. When monitoring is disabled, the member port is added back to the LAG if it has not already been added. This allows a member port to be added back to LAG even though connectivity to the host is down.

Example

The following command disables port 12:

```
configure sharing health-check member-port 12 disable tcp-tracking
```

History

This command was first available in ExtremeXOS 12.1.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure sharing port lacp activity-mode

```
configure sharing port lacp activity-mode [active | passive]
```

Description

Configures whether the switch sends LACPDUs periodically (active) or only in response to LACPDUs sent from the partner on the link (passive).

Syntax Description

<i>port</i>	Specifies the master logical port for the <i>LAG</i> you are setting the activity mode for.
active	Enter this value to have the switch periodically sent LACPDUs for this LAG.
passive	Enter this value to have the switch only respond to LACPDUs for this LAG.

Default

Active.

Usage Guidelines

You must enable sharing and create the LAG prior to assigning this LACP activity mode.



Note

One side of the link must be in active mode in order to pass traffic. If you configure your side in the passive mode, ensure that the partner link is in LACP active mode.

To verify the LACP activity mode, use the `show lacp lag group-id detail` command.

If you attempt to enter a port number that is different than a LAG group ID, the system returns the following error message:

```
ERROR: LAG group Id does not exist
```



Note

In ExtremeXOS version 11.3, the activity mode cannot be changed from active.

Example

The following command changes the activity mode to passive for the specified LAG group ID:

```
configure sharing 5:1 lacp activity-mode passive
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure sharing lacp defaulted-state-action

```
configure sharing port lacp defaulted-state-action [add | delete]
```

Description

Configures a defaulted LAG port to be removed from the aggregator.

Syntax Description

<i>port</i>	Specifies the master logical port for the LAG you are setting the default action for.
add	Enter this value to have the switch add defaulted ports to the aggregator for this LAG.
delete	Enter this value to have the switch delete defaulted ports from the aggregator for this LAG.

Default

Delete.

Usage Guidelines

You must enable sharing and create the LAG prior to configuring this LACP parameter.

You can configure whether you want a defaulted LAG port removed from the aggregator or added back into the aggregator. If you configure the LAG to remove ports that move into the default state, those ports are removed from the aggregator and the port state is set to unselected.



Note

In ExtremeXOS version 11.3, defaulted ports in the LAG are always removed from the aggregator; this is not configurable.

If you configure the LAG to add the defaulted port into the aggregator, the system takes inventory of the number of ports currently in the aggregator:

- If there are fewer ports in the aggregator than the maximum number allowed, the system adds the defaulted port to the aggregator (port set to selected and collecting-distributing).
- If the aggregator has the maximum ports, the system adds the defaulted port to the standby list (port set to standby).

**Note**

If the defaulted port is assigned to standby, that port automatically has a lower priority than any other port in the LAG (including those already in standby).

To verify the LACP default action, use the `show lacp lag group-id detail` command.

If you attempt to enter a port number that is different than a LAG group ID, the system returns the following error message:

```
ERROR: LAG group Id does not exist
```

**Note**

To force the LACP trunk to behave like a static sharing trunk, use this command to add ports to the aggregator.

Example

The following command deletes defaulted ports from the aggregator for the specified LAG group ID:

```
configure sharing 5:1 lacp defaulted-state-action delete
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure sharing lacp fallback

```
configure sharing port lacp fallback [enable | disable]
```

Description

This command provides the ability to configure fallback. If fallback is enabled and LACP PDUs are not received on LACP-configured ports within the timeout period, the port with the lowest priority value will be added to the aggregator. The port stays in this state until fallback is disabled or until LACP PDUs are exchanged between the switch and its link partner, causing LAG reconfiguration.

Syntax Description

<i>port</i>	LAG group ID.
<i>fallback</i>	Allow a single member port with lowest value priority to be added to the aggregator is LACP PDUs are not received within timeout.
enable	Enable fallback. Port priority and fallback timeout control port aggregator membership.
disable	Disable fallback. LACP PDUs or defaulted-state-action control port aggregator membership.

Default

Disabled.

Example

```
* # show lacp lag 17 detail
Lag   Actor   Actor   Partner           Partner   Partner   Agg   Actor
      Sys-Pri Key    MAC              Sys-Pri   Key      Count MAC
-----
17      0    0x03f9 00:00:00:00:00:00      0    0x0000      1 00:04:96:6d:55:13
Enabled      : Yes
LAG State      : Up
Unack count    : 0
Wait-for-count : 0
Current timeout : Long
Activity mode   : Active
Defaulted Action : Delete
Fallback        : Enabled
Fallback timeout : 40 seconds
Receive state   : Enabled
Transmit state  : Enabled
Minimum active  : 1
Selected count  : 1
Standby count   : 0
LAG Id flag     : Yes
  S.pri:0      , S.id:00:04:96:6d:55:13, K:0x03f9
  T.pri:0      , T.id:00:00:00:00:00:00, L:0x0000

Port list:
Member  Port    Rx      Sel      Mux      Actor      Partner
Port    Priority State    Logic    State     Flags      Port
-----
17      10      Initialize Unselected Detached  A-G----- 0
18      5      Initialize Fallback Collect-Dist A-GSCD-- 1018
19      5      Idle      Unselected Detached  ----- 0
=====
Actor Flags: A-Activity, T-Timeout, G-Aggregation, S-Synchronization
              C-Collecting, D-Distributing, F-Defaulted, E-Expired
```

History

This command was first available in ExtremeXOS 21.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure sharing lacp fallback timeout

```
configure sharing port lacp fallback timeout seconds
```

Description

This command configures the LACP fallback timeout value in seconds.

Syntax Description

lacp	LACP (Link Aggregation Control Protocol).
fallback	Allow single member port with lowest value priority to be added to the aggregator if LACP PDUs are not received within timeout.
timeout	Timeout used to determine how long to wait for LACP PDUs before entering fallback.
seconds	Fallback timeout in seconds. Range 0-100.

Default

60 seconds.

Example

```
* # show lacp lag 17 detail
Lag Actor Actor Partner Partner Partner Agg Actor
    Sys-Pri Key MAC Sys-Pri Key Count MAC
-----
17      0 0x03f9 00:00:00:00:00:00      0 0x0000      1 00:04:96:6d:55:13
Enabled      : Yes
LAG State    : Up
Unack count  : 0
Wait-for-count : 0
Current timeout : Long
Activity mode : Active
Defaulted Action : Delete
Fallback     : Enabled
Fallback timeout : 40 seconds
Receive state : Enabled
Transmit state : Enabled
Minimum active : 1
Selected count : 1
Standby count : 0
LAG Id flag   : Yes
S.pri:0      , S.id:00:04:96:6d:55:13, K:0x03f9
T.pri:0      , T.id:00:00:00:00:00:00, L:0x0000

Port list:
Member Port Rx Sel Mux Actor Partner
Port Priority State Logic State Flags Port
```



```

-----
17          10      Initialize  Unselected  Detached    A-G-----  0
18          5       Initialize  Fallback    Collect-Dist A-GSCD-- 1018
19          5        Idle       Unselected  Detached    -----  0
=====
Actor Flags: A-Activity, T-Timeout, G-Aggregation, S-Synchronization
              C-Collecting, D-Distributing, F-Defaulted, E-Expired

```

History

This command was first available in ExtremeXOS 21.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure sharing lacp system-priority

```
configure sharing port lacp system-priority priority
```

Description

Configures the system priority used by LACP for each *LAG* to establish the station on which end assumes priority in determining those LAG ports moved to the collecting/distributing state of the protocol. That end of the LAG with the lowest system priority is the one that assumes control of the determination. This is optional; if you do not configure this parameter, LACP uses system MAC values to determine priority. If you choose to configure this parameter, enter a value between 1 and 65535.

Syntax Description

<i>port</i>	Specifies the master logical port for the LAG you are setting the priority for.
<i>priority</i>	Enter the value you want for the priority of the system for the LACP. The range is 0 to 65535; there is no default.

Default

N/A.

Usage Guidelines

The LACP uses the system MAC values to assign priority to one of the systems, and that system then determines which LAG ports move into the collecting/distributing state and exchange traffic. That end of the LAG with the lowest system priority is the one that assumes control of the determination. If you wish to override the default LACP system priority for a specific LAG, use this command to assign that LAG a specific LACP priority. Enter a value between 0 and 65535.

You must enable sharing and create the LAG prior to assigning this LACP priority.

To verify the LACP system priority, use the `show lacp` command.

To change the system priority you previously assigned to a specific LAG, issue the `configure sharing lacp system-priority` command using the new priority you want. To remove the assigned system priority entirely and use the LACP priorities, issue the `configure sharing lacp system-priority` command using a value of 0.

Example

The following command assigns LAG 10 an LACP system priority of 3:

```
configure sharing 10 lacp system-priority 3
```

History

This command was first available in ExtremeXOS 11.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure sharing lacp timeout

```
configure sharing port lacp timeout [long | short]
```

Description

Configures the timeout used by each *LAG* to stop transmitting once LACPDUs are no longer received from the partner link. You can configure this timeout value to be either 90 seconds, long, or 3 seconds, short.

Syntax Description

<i>port</i>	Specifies the master logical port for the LAG you are setting the timeout value for.
long	Enter this value to use 90 seconds as the timeout value.
short	Enter this value to use 3 seconds as the timeout value.

Default

Long.

Usage Guidelines

You must enable sharing and create the LAG prior to assigning this LACP timeout value.

To verify the LACP timeout value, use the `show lacp lag group-id detail` command.

If you attempt to enter a port number that is different than a LAG group ID, the system returns the following error message:

```
ERROR: LAG group Id does not exist
```



Note

In ExtremeXOS version 11.3, the timeout value is set to long and cannot be changed.

Example

The following command changes the timeout value for the specified LAG group ID to short:

```
configure sharing 5:1 lacp timeout short
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure sharing minimum-active

```
configure sharing port minimum-active min_links_active
```

Description

This command allows you to configure a value for the minimum number of active links to keep the entire *LAG* up.

Syntax Description

sharing	Load sharing.
<i>port</i>	Master port.
minimum-active	Minimum active links for group to remain in service.
<i>min_links_active</i>	Number of active links. Default is 1. Range is 1 – 8.

Default

1

Usage Guidelines

Use this command to configure the value for the minimum number of active links to keep the LAG up.

Example

The following example display output from the `show port port sharing` command using minimum active links:

```
# sh ports 14 sharing
Load Sharing Monitor
Config  Current  Agg      Min    Ld Share  Ld Share  Agg    Link  Link Up
Master  Master   Control  Active Algorithm Group      Mbr    State Transitions
=====
      14          Static      2      L2        14      -      R      0
      1          L2        15      Y      A
      L2        16      -      R      0
=====
Link State: A-Active, D-Disabled, R-Ready, NP-Port not present, L-Loopback
Minimum Active: (<) Group is down. # active links less than configured minimum
Load Sharing Algorithm: (L2) Layer 2 address based, (L3) Layer 3 address based
                      (L3_L4) Layer 3 address and Layer 4 port based
                      (custom) User-selected address-based configuration
Custom Algorithm Configuration: ipv4 L3-and-L4, xor
Number of load sharing trunks: 2 (1 displayed)
```

History

This command was first available in ExtremeXOS 15.7.

Platform Availability

All ExtremeXOS-based platforms that support static LAG and LACP are supported.

configure sharing port-based key

```
configure sharing [ load_sharing_key | default] ports port_list
```

Description

Sets the `load_sharing_key` for all ports in the `port_list`.

Syntax Description

<code>load_sharing_key</code>	Specifies the load sharing key. Valid load sharing keys are in the range [0-15].
default	Unconfigures and resets the load sharing keys for ports in the <code>port_list</code> to default values.
ports	Specifies the logical port for a load-sharing group.
<code>port_list</code>	Specifies one or more ports or slots and ports to be grouped in the <code>LAG</code> .

Default

N/A.

Usage Guidelines

This command sets the *load_sharing_key* for all ports in the *port_list*. **default** unconfigures and resets the load sharing keys for ports in *port_list* to default values.

Configured load sharing keys are displayed in the output of the `show configuration hal` command. Both configured and default load sharing keys are displayed in the output of the "show sharing port-based keys" command.

Example

The following example causes all packets received on ports in slot 1 to choose the lowest port number in all aggregators for distribution.:

```
configure sharing port-based key 0 ports 1
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure slot description

```
configure slot slot description [ slot_description | none ]
```

Description

Adds or removes a descriptive name to a slot.

Syntax Description

slot	Specifies a specific port extender by slot number.
<i>slot</i>	Specifies a specific port extender by slot number.
description	Specifies naming the BPE at the designated slot.
<i>slot_description</i>	Name for the BPE at the designated slot (max. of 64 characters long).
none	Specifies removing the current name assigned to the BPE at the designated slot.

Default

N/A.

Usage Guidelines

To remove a name from a slot, use the **none** option.

The slot name can be up to 64 characters long.

To view a slot's name, use any of the following commands:

- `show slot {slot} {detail} | detail`
- `show vpex bpe`
- `show vpex bpe {slot slot_num} {statistics} {detail}`
- `show vpex bpe {slot slot_num} {environment}`

Example

The following example applies the name "Accounting Dept" to the BPE at slot 100:

```
configure slot 100 description Accounting Dept
```

History

This command was first available in ExtremeXOS 22.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure slot module

```
configure slot slot module module_type
```

Description

Configures a slot for a particular type of node.

Syntax Description

<i>slot</i>	Specifies the slot number.
<i>module_type</i>	The particular switch in the SummitStack.

Usage Guidelines

The `configure sharing lacp timeout` command displays different switch parameters depending on the type of switch you are configuring and the version of ExtremeXOS running on the switch.

Upon powering up the stack, ExtremeXOS automatically determines the system power budget and protects the switch from any potential overpower configurations. If power is available, ExtremeXOS powers on and initializes the nodes in the stack. When ExtremeXOS detects that a node will cause an overpower condition, the node remains powered down, and is not initialized. An entry is made to the system log indicating the condition.

The module type must be a switch that supports SummitStack.

Example

The following command configures slot 2 in a stack for a ExtremeSwitching 5520-24T switch:

```
# configure slot 2 module 5520-24T
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure slot restart-limit

```
configure slot slot_number restart-limit num_restarts
```

Description

Configures the number of times a slot can be restarted on a failure before it is shut down.

Syntax Description

<i>slot_number</i>	Specifies the slot number.
<i>num_restarts</i>	Specifies the number of times the slot can be restarted. The range is from 0 to 10,000.

Default

The default is 5.

Usage Guidelines

This command allows you to configure the number of times a slot can be restarted on a failure before it is shut down. If the number of failures exceeds the restart-limit, the module goes into a “Failed” state. If that occurs, use the [disable slot](#) and [enable slot](#) commands to restart the module.

Example

The following command configures slot 2 on the switch to be restarted up to 3 times upon a failure:

```
configure slot 2 restart-limit 3
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available only on SummitStack.

configure slpp guard ethertype

```
configure slpp guard ethertype hex
```

Description

Configures the Ethertype that the Simple Loop Protection Protocol (SLPP) Guard feature uses to identify SLPP PDUs.

Syntax Description

slpp	Specifies configuring SLPP.
guard	Specifies disabling a port as soon as an SLPP PDU is received.
ethertype	Specifies selecting the Ethertype used by PDUs of the SLPP protocol.
<i>hex</i>	Specifies the Ethertype value in hexadecimal [0x0600-0xffff]. The default is 0x8102.

Default

By default, the Ethertype is 0x8102.

Usage Guidelines

SLPP is an application that detects loops in a Split Multi-link Trunking (SMLT) network. SLPP Guard is a complementary feature that helps prevent loops in networks by administratively disabling an edge port if a switch receives an SLPP PDU from an SMLT network.

This command configures the Ethernet type field of the packet that SLPP Guard uses to identify SLPP PDUs.

Example

The following example configures the SLPP Guard Ethertype as 0x8110:

```
# configure slpp guard ethertype 0x8110
```

History

This command was available in ExtremeXOS 30.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure slpp guard recovery-timeout

```
configure slpp guard [ports [port_list | all] recovery-timeout [seconds | none]
```

Description

Configures the recovery timeout period for the Simple Loop Protection Protocol (SLPP) Guard feature.

Syntax Description

slpp	Specifies configuring SLPP.
guard	Specifies disabling a port as soon as an SLPP PDU is received.
ports	Specifies ports on which to configure the recovery timeout period.
<i>port_list</i>	Selects which ports to configure the recovery timeout period for (list separated by a comma or -).
all	Specifies configuring all ports with the designated recovery timeout period.
recovery-timeout	Specifies configuring the timeout period after which ports are re-enabled.
<i>seconds</i>	Designates the recovery timeout period in seconds after which the ports are re-enabled. Range is 10–65,535. Default is 60 seconds.
none	

Default

By default, the recovery timeout period is 60 seconds.

Usage Guidelines

SLPP is an application that detects loops in a Split Multi-link Trunking (SMLT) network. SLPP Guard is a complementary feature that helps prevent loops in networks by administratively disabling an edge port if a switch receives an SLPP PDU from an SMLT network.

On a port with SLPP Guard enabled, if an SLPP PDU is received, the port is immediately disabled. After the configured timeout value set by this command expires (associated with each port), the port is automatically re-enabled.

Example

The following example configures the recovery timeout period to 600 seconds for port 9:

```
# configure slpp guard ports 9 recovery-timeout 600
```

History

This command was available in ExtremeXOS 30.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure snmp access-profile

```
configure snmp access-profile [ access_profile {readonly | readwrite} |
    [[add rule ] [first | [[before | after] previous_rule]] ] | delete
    rule | none ]
```

Description

Configures *SNMP* to use an *ACL* policy or ACL rule for access control.

Syntax Description

<i>access_profile</i>	Specifies an ACL policy.
readonly	Specifies that access granted by the specified policy is read only.
readwrite	Specifies that access granted by the specified policy is read/write.
add	Specifies that an ACL rule is to be added to the SNMP application.
<i>rule</i>	Specifies an ACL rule.
first	Specifies that the new rule is to be added before all other rules.
before	Specifies that the new rule is to be added before a previous rule.
after	Specifies that the new rule is to be added after a previous rule.
<i>previous_rule</i>	Specifies an existing rule in the application.
delete	Specifies that the named rule is to be deleted.
none	Specifies that all the rules or a policy file is to be deleted.

Default

SNMP access is enabled by default, with no ACL policies.

Usage Guidelines

You must be logged in as administrator to configure SNMP parameters. You can restrict SNMP access in the following ways:

- Implement an ACL policy. You create an ACL policy file that permits or denies a specific list of IP addresses and subnet masks for SNMP. You must create the ACL policy file before you can use this command. If the ACL policy file does not exist on the switch, the switch returns an error message indicating that the file does not exist.

In the ACL policy file for SNMP, the source-address field is the only supported match condition. Any other match conditions are ignored.

Use the `none` option to remove a previously configured ACL policy.

- Add an ACL rule to the SNMP application through this command. Once an ACL is associated with SNMP, all the packets that reach an SNMP module are evaluated with this ACL and appropriate action (permit or deny) is taken, as is done using policy files.

The permit or deny counters are also updated accordingly, regardless of whether the ACL is configured to add counters. To display counter statistics, use the `show access-list counters process snmp` command.

Only the following match conditions and actions are copied to the client memory. Others that may be in the rule are not copied.

Match conditions:

- Source-address—IPv4 and IPv6
- Actions:
 - Permit
 - Deny

When adding a new rule, use the `first`, `before`, and `after previous_rule` parameters to position it within the existing rules.

If the SNMP traffic does not match any of the rules, the default behavior is deny.

Creating an ACL Policy File

To create an ACL policy file, use the `edit policy` command. For more information about creating and implementing ACL policy files, see the *Policy Manager* and *ACLs* chapters in the [Switch Engine 32.3 User Guide](#).

If you attempt to implement a policy that does not exist, an error message similar to the following appears:

```
Error: Policy /config/MyAccessProfile.pol does not exist on file system
```

If this occurs, make sure the policy you want to implement exists. To confirm the existence of the policies, use the `ls` command. If the policy does not exist, create the ACL policy file.

Viewing SNMP Information

To display the current management configuration, including SNMP access related information, whether SNMP access is enabled or disabled, and whether any ACL or rules are configured for SNMP, use the following command: `show management`

Example

The following example applies the ACL policy file `MyAccessProfile_2` to SNMP:

```
configure snmp access-profile MyAccessProfile_2
```

The following example applies the ACL rule DenyAccess to SNMP as the first rule in the list:

```
configure snmp access-profile add DenyAccess first
```

The following example deletes the ACL rule DenyAccess from the SNMP application:

```
configure snmp access-profile delete DenyAccess
```

To delete the use of all the ACL rules or a policy file by SNMP, use the following command:

```
configure snmp access-profile none
```

History

This command was first available in ExtremeXOS 11.6.

Support for individual ACL rules was added in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure snmp add community

```
configure snmp add community [readonly | readwrite] alphanumeric_string
    [encrypted enc_community_name | community name | hex
    hex_community_name ] store-encrypted
```

Description

Adds a *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	Community name is encrypted.
hex	Provide value in hexadecimal.
<i>hex_community_name</i>	Community name in hexadecimal.
store-encrypted	Community name will be stored as encrypted, instead plain text.
<i>alphanumeric_string</i>	Specifies an SNMP community string name. See “Usage Guidelines” for more information.

Default

N/A

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. Sixteen read-only and sixteen read/write community strings can be configured on the switch, including the defaults.

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 `configure snmp add community` command allows you to add multiple community strings in addition to the default community string.

An SNMP community string can contain up to 32 characters.

We recommend that you delete the defaults of the community strings. To delete the value of the default read/write and read-only community strings, use the `configure snmp delete community` command.

Example

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

```
configure snmp add community readonly hex 65:01
```

History

This command was first available in ExtremeXOS 10.1.

The **hex** keyword and `hex_community_name` variable were added in ExtremeXOS 15.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure snmp add notification-log

```
configure snmp add notification-log [default | name | hex hex_name ]  
    user [snmp_user_name | hex hex_snmp_user_name ] sec-model sec_model  
    sec-level sec_level ]
```

Description

Adds a notification log.

Syntax Description

<i>name</i>	Specifies the name of the log.
hex	Provide value in hexadecimal.
<i>hex_name</i>	Name of the log in hexadecimal.
user	Name of the <i>SNMP</i> user on whose behalf the log should be created.
<i>snmp_user_name</i>	SNMP user name in ASCII..
<i>hex_snmp_user_name</i>	SNMP user name in hexadecimal.
sec-model	Security framework associated with the user.
<i>sec_model</i>	Security model.
sec-level	Authentication and privacy levels of the user.
<i>sec_level</i>	Security level.

Default

Disabled.

Usage Guidelines

Use this command to add a notification log. All entries in the log and its configuration are removed when this command is successfully executed.

Example

The following example adds *nmslog1*:

```
configure snmp add notification log nmslog1 user admin sec-model usm sec-level priv
```

History

This command was first available in ExtremeXOS 15.5.

The **hex** keyword, *hex_name* variable, and *hex_snmp_user_name* variable were added in ExtremeXOS 15.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure snmp add trapreceiver

```
configure snmp add trapreceiver [ip_address | ipv6_address] community
    [[hex hex_community_name] | community_name] {port port_number} {from
    [src_ip_address | src_ipv6_address]} {vr vr_name} {mode trap_mode}
```

Description

Adds the IP address of a trap receiver to the trap receiver list and specifies which SNMPv1/v2c traps are to be sent.

Syntax Description

<i>ip_address</i>	Specifies an <i>SNMP</i> trap receiver IPv4 address.
<i>ipv6_address</i>	Specifies an SNMP trap receiver IPv6 address
<i>hex_community_name</i>	Specifies that the trap receiver is to be supplied as a colon separated string of hex octets.
<i>community_name</i>	Specifies the community string of the trap receiver to be supplied in ASCII format.
<i>port_number</i>	Specifies a UDP port to which the trap should be sent. Default is 162.
<i>src_ip_address</i>	Specifies the IPv4 address of a <i>VLAN</i> to be used as the source address for the trap.
<i>src_ipv6_address</i>	Specifies the IPv6 address of a VLAN to be used as the source address for the trap.
<i>vr_name</i>	Specifies the name of the virtual router.
<i>trap_mode</i>	Specifies the mode of the traps:enhanced—Contains extra varbinds at the end.standard—Does not contain extra varbinds.

Default

Trap receivers are in enhanced mode by default, and the version is SNMPv2c 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 configured to receive the specific trap group.

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.

Example

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

```
configure 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:

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

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:

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

History

This command was first available in ExtremeXOS 10.1.

The virtual router parameter was added in ExtremeXOS 12.3.

IPv6 support was added in ExtremeXOS 12.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure snmp delete community

```
configure snmp delete community [readonly | readwrite] [all |  
    community_name | alphanumeric_string | hex hex_community_name |  
    encrypted enc_community_name ]
```

Description

Deletes a *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.
all	Specifies all of the SNMP community strings.
<i>alphanumeric_string</i>	Specifies an SNMP community string name. See “Usage Guidelines” for more information.
hex	Provide value in hexadecimal.
<i>hex_community_name</i>	Community name in hexadecimal.

Default

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

Usage Guidelines

You must have at least one community string for SNMP access. If you delete all of the community strings on your system, you will no longer have SNMP access, even if you have SNMP enabled.

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. Sixteen read-only and sixteen read-write community strings can be configured on the switch, including the defaults. 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 32 characters.

For increased security, we recommend that you change the defaults of the read/write and read-only community strings.

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

Example

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

```
configure snmp delete community readonly hex 65:01
```

History

This command was first available in ExtremeXOS 10.1.

The **hex** keyword and `hex_community_name` variable were added in ExtremeXOS 15.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure snmp delete notification-log

```
configure snmp delete notification-log [default | name | hex hex_name ]
```

Description

Deletes a notification log.

Syntax Description

default	The default log.
<i>name</i>	Specifies the name of the log.

hex	Provide value in hexadecimal.
<i>hex_name</i>	Name of the log in hexadecimal.

Default

Disabled.

Usage Guidelines

Use this command to delete a notification log. All entries in the log and its configuration are removed when this command is successfully executed.

Example

The following example deletes *nmslog1*:

```
configure snmp delete notification-log hex 01:02
```

History

This command was first available in ExtremeXOS 15.5.

The **default** and **hex** keywords and *hex_name* variable were added in ExtremeXOS 15.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure snmp delete trapreceiver

```
configure snmp delete trapreceiver [[ip_address | ipv6_address]  
{port_number} | all]
```

Description

Deletes a specified trap receiver or all authorized trap receivers.

Syntax Description

<i>ip_address</i>	Specifies an <i>SNMP</i> trap receiver IPv4 address.
<i>ipv6_address</i>	Specifies an SNMP trap receiver IPv6 address.
<i>port_number</i>	Specifies the port associated with the receiver.
all	Specifies all SNMP trap receiver IP addresses.

Default

The default port number is 162.

Usage Guidelines

Use this command to delete a trap receiver of the specified IPv4 or IPv6 address, or all authorized trap receivers.

This command deletes only the first SNMPv1/v2c trap receiver whose IP address and port number match the specified value.

Example

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

```
configure snmp delete trapreceiver 10.101.0.100
```

The following command deletes entries in the trap receiver list for 10.101.0.100, port 9990:

```
configure snmp delete trapreceiver 10.101.0.100 9990
```

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

History

This command was first available in ExtremeXOS 10.1.

IPv6 support was added in ExtremeXOS 12.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure snmp ifmibifalias size

```
config snmp ifmib ifalias size [default | extended ]
```

Description

Controls the accessible string size for the *SNMP* ifAlias object.

Syntax Description

default	Specifies read-only access to the system.
extended	Specifies read and write access to the system.

Default

N/A.

Usage Guidelines

Use this command to control the accessible string size for the SNMP ifAlias object.

If you choose the extended size option, the following warning will be displayed:

Warning: Changing the size to [extended] requires the use of increased 255 chars long ifAlias object of ifXtable from IF-MIB(RFC 2233)

You can always configure a 255 character long string regardless the configured value of ifAlias size. Its value only affects the SNMP behavior.

Example

The following example shows how to configure the accessible string size for the SNMP ifAlias to the default value:

```
config snmp ifmib ifalias size[default
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure snmp notification-log filter-profile-name

```
configure snmp notification-log [ default | name | hex
    hex_name ] [ filter-profile-name [ none | filter_profile_name
    | hex hex_filter_profile_name ] | entry-limit [system-managed |
    entry_limit ] ]
```

Description

Changes the configuration of a notification log.

Syntax Description

filter-profile-name	Sets the notification filter profile for this log (default: none).
none	Specifies the global entry limit.
<i>filter_profile_name</i>	Sets the notification filter profile for this log (default: none)
none	Specifies no notification filter profile.

<i>filter_profile_name</i>	Specifies the notification filter profile name.
hex	Provide value in hexadecimal.
<i>hex_name</i>	Name of the log in hexadecimal.
filter-profile-name	Set the notification filter profile for this log. Default = none.
<i>hex_filter_profile_name</i>	Notification filter profile name in hexadecimal.
entry-limit	Sets the maximum number of entries in this log (default: system-managed).
system-managed	Specifies that entries will be removed from this log when the global entry limit is exceeded.
<i>entry_limit</i>	Specifies that entries will be removed from this log when the specified limit is exceeded. The range is 1-16000.

Default

Usage Guidelines

Use this command to change the configuration of a notification log. Use the `configure snmpv3 add filter-profile` command to create notification filter profiles.

Example

The following example sets the filter for the default log to all and its maximum size to 1500:

```
configure snmp notification-log default filter-profile-name all entry-limit 1500
```

History

This command was first available in ExtremeXOS 15.5.

The **hex** keywords and *hex_name* and *hex_filter_profile_name* variables were added in ExtremeXOS 15.6

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure snmp notification-log

```
configure snmp notification-log [global-entry-limit global_entry_limit |  
global-age-out [none | minutes]]
```

Description

Configures notification log settings that affect all logs.

Syntax Description

global-entry-limit	Sets the maximum number of entries in all logs combined (default: 16000)
<i>global_entry_limit</i>	Specifies the global entry limit. Range is 1-16000.
global-age-out	Sets the number of minutes a notification should be kept in a log before it is automatically removed (default: 1440 for one day).
none	Specifies that notifications are not aged out.
<i>minutes</i>	Specifies the global age out in minutes. The range is 1 - 4294967295.

Default

global-entry-limit is 16000.

global-age-out is 1440 for one day.

Usage Guidelines

Use this command to configure notification log settings that affect all logs.

Example

The following example sets the log size to 10000, and disable aging:

```
configure snmp notification-log global-entry-limit 10000 global-age-out none
```

History

This command was first available in ExtremeXOS 15.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure snmp sysContact

```
configure snmp syscontact sysContact
```

Description

Configures the name of the system contact.

Syntax Description

<i>sysContact</i>	An alphanumeric string that 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. The allowed character set is A-Z, a-z, 0-9, +, -, @, _, :, ;, (,) / " .

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.

Example

The following example defines FredJ as the system contact:

```
configure snmp syscontact FredJ
```

The following output from the `show switch` command displays FredJ as the system contact:

```
SysName:      engineeringlab
SysLocation:  englab
SysContact:   FredJ
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure snmp sysLocation

```
configure snmp syslocation sysLocation
```

Description

Configures the location of the switch.

Syntax Description

<i>sysLocation</i>	An alphanumeric string that 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. The allowed character set is A-Z, a-z, 0-9, +-@_.,;()/ " .

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.

Example

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

```
configure snmp syslocation englab
```

The following output from the `show switch` command displays englab as the location of the switch:

```
SysName:      engineeringlab
SysLocation:  englab
SysContact:   FredJ
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure snmp sysName

```
configure snmp sysname sysName
```

Description

Configures the name of the switch.

Syntax Description

<i>sysName</i>	An alphanumeric string that specifies a device name.
----------------	--

Default

The default sysName is the model name of the device (for example, ExtremeSwitching X440-G2).

Usage Guidelines

You can use this command to change the name of the switch. A maximum of 255 characters is allowed. The allowed character set is A-Z, a-z, 0-9, +-@_.,;()/ " .

The sysName appears in the switch prompt. On a SummitStack, the sysName appears in the prompt of all active nodes in the stack when there is a master node present in the stack.

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.

Example

The following example names the switch:

```
configure snmp sysname engineeringlab
```

The following output from the `show switch` command displays engineeringlab as the name of the switch:

```
SysName:      engineeringlab
SysLocation:  englab
SysContact:   FredJ
```

History

This command was first available in ExtremeXOS 10.1.

Beginning in ExtremeXOS 15.7, the maximum number of characters has been changed to 255.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure snmp traps batch-delay bfd

```
configure snmp traps batch-delay bfd none | delay
```

Description

This command allows you to configure the time during which the set of affected sessions will be collected and a single trap will be set for contiguous session IDs. This means that there is a small delay between event occurrence and trap generation. You have the option to disable this optimization delay using the **none** option.

Syntax Description

snmp	Configure <i>SNMP</i> specific settings.
traps	Configure SNMP Trap generation settings.
batch-delay	Maximum delay before trap generation in order to combine multiple traps into a single trap.

none	Disables trap optimization which results in generation of one trap for status change of each session.
<i>delay</i>	Choose delay to balance between number of traps and delay in trap generation. Range is 50 to 65535 ms.

Default

1000 ms.

Usage Guidelines

Use this command to configure the time window during which the set of affected sessions is collected and single trap is set for contiguous sessions IDs.

Example

The following command configures the BFD batch-delay:

```
# configure snmp traps batch-delay bfd 1000
```

History

This command was first available in ExtremeXOS 15.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure snmpv3 add access

```
configure snmpv3 add access [[hex hex_group_name] | group_name] {sec-  
model [snmpv1 | snmpv2c | usm]} {sec-level [noauth | authnopriv  
| priv]} {read-view [[hex hex_read_view_name] | read_view_name]}  
{write-view [[hex hex_write_view_name] | write_view_name]} {notify-  
view [[hex hex_notify_view_nam]] | notify_view_name]} {volatile}
```

Description

Creates (and modifies) a group and its access rights.

Syntax Description

<i>hex_group_name</i>	Specifies the group name to add or modify. The value is to be supplied as a colon separated string of hex octets.
<i>group_name</i>	Specifies the group name to add or modify. The value is to be supplied in ASCII format.
sec-model	Specifies the security model to use.

snmpv1	Specifies the SNMPv1 security model.
snmpv2c	Specifies the SNMPv2c security model.
usm	Specifies the SNMPv3 User-based Security Model (USM).
sec-level	Specifies the security level for the group.
noauth	Specifies no authentication (and implies no privacy) for the security level.
authnopriv	Specifies authentication and no privacy for the security level.
priv	Specifies authentication and privacy for the security level.
read-view	Specifies the read view name:hex_read_view_name—Specifies a hex value supplied as a colon separated string of hex octetsread_view_name—Specifies an ASCII value.
write-view	Specifies the write view name:hex_write_view_name—Specifies a hex value supplied as a colon separated string of hex octetswrite_view_name—Specifies an ASCII value.
notify-view	Specifies the notify view name:hex_notify_view_name—Specifies a hex value supplied as a colon separated string of hex octetsnotify_view_name—Specifies an ASCII value.
volatile	Specifies volatile storage.

Default

The default values are:

- sec-model—USM
- sec-level—noauth
- read view name—defaultUserView
- write view name— ""
- notify view name—defaultNotifyView
- non-volatile storage

Usage Guidelines

Use this command to configure access rights for a group. All access groups are created with a unique default context, "", as that is the only supported context.

Use more than one character when creating unique community strings and access group names.

A number of default groups are already defined. These groups are: admin, initial, v1v2c_ro, v1v2c_rw.

- The default groups defined are v1v2c_ro for security name v1v2c_ro, v1v2c_rw for security name v1v2c_rw, admin for security name admin, and initial for security names initial, initialmd5, initialsha, initialmd5Priv and initialshaPriv.
- The default access defined are admin, initial, v1v2c_ro, v1v2c_rw, and v1v2cNotifyGroup.

Example

In the following command, access for the group defaultROGroup is created with all the default values: security model usm, security level noauth, read view defaultUserView, no write view, notify view defaultNotifyView, and storage nonvolatile.

```
configure snmpv3 add access defaultROGroup
```

In the following command, access for the group defaultROGroup is created with the values: security model USM, security level authnopriv, read view defaultAdminView, write view defaultAdminView, notify view defaultAdminView, and storage nonvolatile.

```
configure snmpv3 add access defaultROGroup sec-model usm sec-level authnopriv read-view
defaultAdminView write-view defaultAdminView notify-view defaultAdminView
```

History

This command was first available in ExtremeXOS 10.1.

The hex_read_view_name, hex_write_view_name, and hex_notify_view_name parameters were added in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure snmpv3 add community

```
configure snmpv3 add community [[hex hex_community_index] |
community_index] [encrypted name community_name | name [[hex
hex_community_name] | community_name] {store-encrypted} ] user
[[hex hex_user_name] | user_name] {tag [[hex transport_tag] |
transport_tag]} {volatile}
```

Description

Adds an SNMPv3 community entry.

Syntax Description

<i>hex_community_index</i>	Specifies the row index in the snmpCommunity table as a hex value supplied as a colon separated string of hex octets.
<i>community_index</i>	Specifies the row index in the snmpCommunity Table as an ASCII value.
<i>hex_community_name</i>	Specifies the community name as a hex value supplied as a colon separated string of hex octets.
<i>community_name</i>	Specifies the community name as an ASCII value.

<i>hex_user_name</i>	Specifies the USM user name as a hex value supplied as a colon separated string of hex octets.
<i>user_name</i>	Specifies the USM user name as an ASCII value.
tag	Specifies the tag used to locate transport endpoints in SnmpTargetAddrTable. When this community entry is used to authenticate v1/v2c messages, this tag is used to verify the authenticity of the remote entity. <i>hex_transport_tag</i> —Specifies a hex value supplied as a colon separated string of hex octets <i>transport_tag</i> —Specifies an ASCII value
volatile	Specifies volatile storage.

Default

N/A.

Usage Guidelines

Use this command to create or modify an SMMPv3 community in the community MIB.

Example

```
switch # configure snmp add community readonly extreme store-encrypted
switch # show snmpv3 community
Community Index : extreme
Community Name : hys{fnj (encrypted)
Security Name : vlv2c_ro
Context EngineID : 80:00:07:7c:03:00:04:96:27:b6:63
Context Name :
Transport Tag :
Storage Type : NonVolatile
Row Status : Active
switch # configure snmp add community readwrite extreme123
switch # show snmpv3 community
Community Index : extreme
Community Name : hys{fnj (encrypted)
Security Name : vlv2c_ro
Context EngineID : 80:00:07:7c:03:00:04:96:27:b6:63
Context Name :
Transport Tag :
Storage Type : NonVolatile
Row Status : Active
Community Index : extreme123
Community Name : extreme123
Security Name : vlv2c_rw
Context EngineID : 80:00:07:7c:03:00:04:96:27:b6:63
Context Name :
Transport Tag :
Storage Type : NonVolatile
Row Status : Active
switch # show configuration "snmp"
#
# Module snmpMaster configuration.
#
configure snmpv3 add community extreme encrypted name hys{fnj user vlv2c_ro
configure snmpv3 add community extreme123 name extreme123 user vlv2c_rw
The following command creates an entry with the community index comm_index, community
```

```
name comm_public, and user (security) name vlv2c_user:
configure snmpv3 add community comm_index name comm_public user vlv2c_user
```

History

This command was first available in ExtremeXOS 10.1.

The hex_community_index, hex_community_name, hex_user_name, and hex_transport_tag parameters were added in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure snmpv3 add filter

```
configure snmpv3 add filter [[hex hex_profile_name] | profile_name]
    subtree object_identifier {/subtree_mask} type [included | excluded]
    {volatile}
```

Description

Adds a filter to a filter profile.

Syntax Description

<i>hex_profile_name</i>	Specifies the filter profile that the current filter is added to. The value is to be supplied as a colon separated string of hex octets.
<i>profile_name</i>	Specifies the filter profile that the current filter is added to in ASCII format.
<i>object_identifier</i>	Specifies a MIB subtree.
subtree_mask	Specifies a hex octet string used to mask the subtree. For example, f7a indicates 1.1.1.0.1.1.1.0.1.0.
included	Specifies that the MIB subtree defined by object identifier/mask is to be included.
excluded	Specifies that the MIB subtree defined by object identifier/mask is to be excluded.
volatile	Specifies volatile storage.

Default

The default values are:

- mask value—empty string (all 1s).
- type—included.
- storage—non-volatile.

Usage Guidelines

Use this command to create a filter entry in the snmpNotifyFilterTable. Each filter includes or excludes a portion of the MIB. Multiple filter entries comprise a filter profile that can eventually be associated with a target address. Other commands are used to associate a filter profile with a parameter name, and the parameter name with a target address.

This command can be used multiple times to configure the exact filter profile desired.

Example

The following command adds a filter to the filter profile prof1 that includes the MIB subtree 1.3.6.1.4.1/f0:

```
configure snmpv3 add filter prof1 subtree 1.3.6.1.4.1/f0 type included
```

History

This command was first available in ExtremeXOS 10.1.

The hex_profile_name parameter was added in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure snmpv3 add filter-profile

```
configure snmpv3 add filter-profile [[hex hex_profile_name] |  
profile_name] param [[hex hex_param_name]] | param_name] {volatile}
```

Description

Associates a filter profile with a parameter name.

Syntax Description

<i>hex_profile_name</i>	Specifies the filter profile name. The value is to be supplied as a colon separated string of hex octets.
<i>profile_name</i>	Specifies the filter profile name in ASCII format.
<i>hex_param_name</i>	Specifies a parameter name to associate with the filter profile. The value to follow is to be supplies as a colon separated string of hex octets.
<i>param_name</i>	Specifies a parameter name to associate with the filter profile in ASCII format.
volatile	Specifies volatile storage.

Default

The default storage type is non-volatile.

Usage Guidelines

Use this command to add an entry to the snmpNotifyFilterProfileTable. This table associates a filter profile with a parameter name. The parameter name is associated with target addresses, and the filter profile is associated with a series of filters, so, in effect, you are associating a series of filters with a target address.

Example

The following command associates the filter profile prof1 with the parameter name P1:

```
configure snmpv3 add filter-profile prof1 param P1
```

History

This command was first available in ExtremeXOS 10.1.

The hex_profile_name and hex_param_name parameters were added in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure snmpv3 add group user

```
configure snmpv3 add group [[hex hex_group_name] | group_name] user
    [[hex hex_user_name] | user_name] {sec-model [snmpv1 | snmpv2c | usm]}
    {volatile}
```

Description

Adds a user name (security name) to a group.

Syntax Description

<i>hex_group_name</i>	Specifies the group name to add or modify. The value is to be supplied as a colon separated string of hex octets.
<i>group_name</i>	Specifies the group name to add or modify in ASCII format.
<i>hex_user_name</i>	Specifies the user name to add or modify. The value to follow is to be supplies as a colon separated string of hex octets.
<i>user_name</i>	Specifies the user name to add or modify in ASCII format.
sec-model	Specifies the security model to use.
snmpv1	Specifies the SNMPv1 security model.

snmpv2c	Specifies the SNMPv2c security model.
usm	Specifies the SNMPv3 User-based Security Model (USM).
volatile	Specifies volatile storage.

Default

The default values are:

- sec-model—USM.
- non-volatile storage.

Usage Guidelines

Use this command to associate a user name with a group.

As per the SNMPv3 RFC, a security name is model independent while a username is model dependent. For simplicity, both are assumed to be same here. User names and security names are handled the same. In other words, if a user is created with the user name username, the security name value is the same, username.

Every group is uniquely identified by a security name and security model. So the same security name can be associated to a group name but with different security models.

Example

The following command associates the user userV1 to the group defaultRoGroup with SNMPv1 security:

```
configure snmpv3 add group defaultRoGroup user userV1 sec-model snmpv1
```

The following command associates the user userv3 with security model USM and storage type volatile to the access group defaultRoGroup:

```
configure snmpv3 add group defaultRoGroup user userV3 volatile
```

History

This command was first available in ExtremeXOS 10.1.

The hex_group_name and hex_user_name parameters were added in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure snmpv3 add mib-view

```
configure snmpv3 add mib-view [[hex hex_view_name] | view_name]
    subtree object_identifier {subtree_mask} {type [included | excluded]}
    {volatile}
```

Description

Adds (and modifies) a MIB view.

Syntax Description

<i>hex_view_name</i>	Specifies the MIB view name to add or modify. The value is to be supplies as a colon separated string of hex octets.
<i>view_name</i>	Specifies the MIB view name to add or modify in ASCII format.
<i>object_identifier</i>	Specifies a MIB subtree.
<i>subtree_mask</i>	Specifies a hex octet string used to mask the subtree. For example, f7a indicates 1.1.1.0.1.1.1.0.1.0.
included	Specifies that the MIB subtree defined by subtree/mask is to be included.
excluded	Specifies that the MIB subtree defined by subtree/mask is to be excluded.
volatile	Specifies volatile storage.

Default

The default mask value is an empty string (all 1s). The other default values are included and non-volatile.

Usage Guidelines

Use this command to create a MIB view into a subtree of the MIB. If the view already exists, this command modifies the view to additionally include or exclude the specified subtree.

In addition to the created MIB views, there are three default views. They are: defaultUserView, defaultAdminView, and defaultNotifyView.

Example

The following command creates the MIB view allMIB with the subtree 1.3 included as non-volatile:

```
configure snmpv3 add mib-view allMIB subtree 1.3
```

The following command creates the view extremeMib with the subtree 1.3.6.1.4.1.1916 included as non-volatile:

```
configure snmpv3 add mib-view extremeMib subtree 1.3.6.1.4.1.1916
```

The following command creates a view vrrpTrapNewMaster which excludes *VRRP* notification .1 and the entry is volatile:

```
configure snmpv3 add mib-view vrrpTrapNewMaster 1.3.6.1.2.1.68.0.1/ff8 type excluded
volatile
```

History

This command was first available in ExtremeXOS 10.1.

The hex_view_name parameter was added in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure snmpv3 add notify

```
configure snmpv3 add notify [[hex hex_notify_name] | notify_name] tag
[[hex hex_tag] | tag] {type [trap | inform]}{volatile}
```

Description

Adds an entry to the snmpNotifyTable.

Syntax Description

<i>hex_notify_name</i>	Specifies the notify name to add. The value is to be supplied as a colon separated string of hex octets.
<i>notify_name</i>	Specifies the notify name to add in ASCII format.
<i>hex_tag</i>	Specifies a string identifier for the notifications to be sent to the target. The value is supplied as a colon separated string of octets.
<i>tag</i>	Specifies a string identifier for the notifications to be sent to the target in ASCII format.
trap	Specifies an unconfirmed notification.
inform	Specifies a confirmed notification.
volatile	Specifies volatile storage. By specifying volatile storage, the configuration is not saved across a switch reboot.

Default

The default storage type is non-volatile.

The default type is trap.

Usage Guidelines

Use this command to add an entry to the snmpNotifyTable. When a notification is to be sent, this table is examined. For the target addresses that have been associated with the tags present in the table, notifications are sent based on the filters also associated with the target addresses.

Example

The following command sends notifications to addresses associated with the tag type1:

```
configure snmpv3 add notify N1 tag type1
```

History

This command was first available in ExtremeXOS 10.1.

The hex_notify_name and hex_tag parameters were added in ExtremeXOS 11.0.

The INFORM option was added in ExtremeXOS 12.5.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure snmpv3 add target-addr

```
configure snmpv3 add target-addr [[hex hex_addr_name] | addr_name]
    param [[hex hex_param_name] |param_name ] ipaddress [ ip_address
    | ipv4-with-mask ip_and_tmask ] | [ ipv6_address | ipv6-with-mask
    ipv6_and_tmask ]] {transport-port port_number} {from [src_ip_address
    | src_ipv6_address]} {vr vr_name} {tag-list [tag_list | hex
    hex_tag_list]} {volatile}
```

Description

Adds and configures an SNMPv3 target address and associates filtering, security, and notifications with that address.

Syntax Description

<i>hex_addr_name</i>	Specifies a string identifier for the target address. The value is to be supplied as a colon separated string of hex octets.
<i>addr_name</i>	Specifies a string identifier for the target address in ASCII format.
<i>hex_param_name</i>	Specifies the parameter name associated with the target. The value is to be supplied as a colon separated string of hex octets.
<i>param_name</i>	Specifies the parameter name associated with the target in ASCII format.
<i>ip_address</i>	Specifies an SNMPv3 target IPv4 address.

ipv4-with-mask	Specify IPv4 address with hexadecimal mask.
<i>ip_and_tmask</i>	Specifies the IPv4 address and hexadecimal mask in form A.B.C.D/NN...
<i>ipv6_address</i>	Specifies an SNMPv3 target IPv6 address.
ipv6-with-mask	Specify IPv6 address with hexadecimal mask.
<i>ipv6_and_tmask</i>	Specifies an IPv6 address and hexadecimal mask in form A:B:C:D:E:F:G:H/NN...
<i>port_number</i>	Specifies a UDP port. Default is 162.
<i>src_ip_address</i>	Specifies the IPv4 address of a VLAN to be used as the source address for the trap.
<i>src_ipv6_address</i>	Specifies the IPv6 address of a VLAN to be used as the source address for the trap.
<i>vr_name</i>	Specifies the name of the virtual router.
tag-list	Specifies a list of comma separated string identifiers for the notifications to be sent to the target.
<i>hex_tag_list</i>	Tag list in RFC 3413 format (in hexadecimal).
volatile	Specifies volatile storage. By specifying volatile storage, the configuration is not saved across a switch reboot.

Default

The default values are:

- transport-port—port 162.
- non-volatile storage.

If you do not specify tag-list the single tag defaultNotify, a pre-defined value in the snmpNotifyTable is used.

Usage Guidelines

Use this command to create an entry in the SNMPv3 snmpTargetAddressTable. The **param** parameter associates the target address with an entry in the snmpTargetParamsTable, which specifies security and storage parameters for messages to the target address, and an entry in the snmpNotifyFilterProfileTable, which specifies filter profiles to use for notifications to the target address. The filter profiles are associated with the filters in the snmpNotifyFilterTable.

The list of tag-lists must match one or more of the tags in the snmpNotifyTable for the trap to be sent out.

Example

The following command specifies a target address of 10.203.0.22 with the name A1, and associates it with the security parameters and target address parameter P1:

```
configure snmpv3 add target-addr A1 param P1 ipaddress 10.203.0.22
```

The following command specifies a target address of 10.203.0.22 with the name A1, and associates it with the security parameters and target address parameter P1, and the notification tags type1 and type2:

```
configure snmpv3 add target-addr A1 param P1 ipaddress 10.203.0.22 from 10.203.0.23 tag-list type1,type2
```

History

This command was first available in ExtremeXOS 10.1.

The virtual router, IP address and hexadecimal mask parameters were added in ExtremeXOS 12.3.

IPv6 support was added in ExtremeXOS 12.4.

The **IPv4-with-mask** and **IPv6-with-mask** keywords were added in ExtremeXOS 15.3.2.

The **hex** keyword and *hex_tag_list* variable were added in ExtremeXOS 15.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure snmpv3 add target-params

```
configure snmpv3 add target-params [[hex hex_param_name] |  
  param_name] user [[hex hex_user_name] | user_name] mp-model [snmpv1  
  | snmpv2c | snmpv3] sec-model [snmpv1 | snmpv2c | usm] {sec-level  
  [noauth | authnopriv | priv]} {volatile}
```

Description

Adds and configures SNMPv3 target parameters.

Syntax Description

<i>hex_param_name</i>	Specifies the parameter name associated with the target. The value is to be supplied as a colon separated string of hex octets.
<i>param_name</i>	Specifies the parameter name associated with the target in ASCII format.
<i>hex_user_name</i>	Specifies a user name. The value is to be supplied as a colon separated string of hex octets.
<i>user_name</i>	Specifies a user name in ASCII format.
mp-model	Specifies a message processing model; choose from SNMPv1, SNMPv2, or SNMPv3.
sec-model	Specifies the security model to use.
snmpv1	Specifies the SNMPv1 security model.

snmpv2c	Specifies the SNMPv2c security model.
usm	Specifies the SNMPv3 User-based Security Model (USM).
sec-level	Specifies the security level for the group.
noauth	Specifies no authentication (and implies no privacy) for the security level.
authnopriv	Specifies authentication and no privacy for the security level.
priv	Specifies authentication and privacy for the security level.
volatile	Specifies volatile storage. By specifying volatile storage, the configuration is not saved across a switch reboot.

Default

The default values are:

- sec-level—noauth.
- non-volatile storage.

Usage Guidelines

Use this command to create an entry in the SNMPv3 snmpTargetParamsTable. This table specifies the message processing model, security level, security model, and the storage parameters for messages to any target addresses associated with a particular parameter name.

To associate a target address with a parameter name, see the command `configure snmpv3 add target-addr`.

Example

The following command specifies a target parameters entry named P1, a user name of guest, message processing and security model of SNMPv2c, and a security level of no authentication:

```
configure snmpv3 add target-params P1 user guest mp-model snmpv2c sec-model snmpv2c sec-level noauth
```

History

This command was first available in ExtremeXOS 10.1.

The hex_param_name and hex_user_name parameters were added in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure snmpv3 add user

```
configure snmpv3 add user [ hex hex_user_name | user_name ]
    {engine-id engine_id} {authentication [md5 | sha] {localized-key
    auth_localized_key | hex hex_auth_password | auth_password} {privacy
    {des | 3des | aes {128 | 192 | 256}} {localized-key priv_localized_key |
    hex hex_priv_password | priv_password} }} {volatile}
```

Description

Adds (and modifies) an SNMPv3 user.

Syntax Description

<i>hex_user_name</i>	Specifies the user name to add or modify. The value is to be supplied as a colon separated string of hex octets.
<i>user_name</i>	Specifies the user name to add or modify in ASCII format.
engine-id	<i>SNMP</i> engine id. If not specified, the user is created with the local engine id.
<i>engine_id</i>	Engine id (in hexadecimal)"; type="ostring_t"
authentication	Specifies the authentication password or hex string to use for generating the authentication key for this user.
md5	Specifies RSA Data Security, Inc. <i>MD5</i> Message-Digest Algorithm authentication.
sha	Specifies SHA authentication.
localized-key	Following value is a MD5 or SHA digest of the engine-id and user's password.
<i>auth_localized_key</i>	Authentication localized key (in hexadecimal));type="ostring_t"
privacy	Specifies the privacy password or hex string to use for generating the privacy key for this user.
des	Specifies the use of the 56-bit DES algorithm for encryption. This is the default.
3des	Specifies the use of the 168-bit 3DES algorithm for encryption.
aes	Specifies the use of the AES algorithm for encryption.
128	Specifies the use of the 128-bit AES algorithm for encryption.
192	Specifies the use of the 192-bit AES algorithm for encryption.
256	Specifies the use of the 256-bit AES algorithm for encryption.
<i>priv_localized_key</i>	Privacy localized key (in hexadecimal)"; type="ostring_t"
volatile	Specifies volatile storage. By specifying volatile storage, the configuration is not saved across a switch reboot.

Default

The default values are:

- authentication—no authentication.
- privacy—no privacy.
- non-volatile storage.

Usage Guidelines

Use this command to create or modify an SNMPv3 user configuration.

The default user names are: admin, initial, initialmd5, initialsha, initialmd5Priv, initialshaPriv. The initial password for admin is password. For the other default users, the initial password is the user name.

If hex is specified, supply a 16 octet hex string for RSA Data Security, Inc. MD5 Message-Digest Algorithm, or a 20 octet hex string for SHA.

You must specify authentication if you want to specify privacy. There is no support for privacy without authentication.



Note

3DES, AES 192, and AES 256 bit encryptions are proprietary implementations and may not work with some SNMP managers.

SNMPv3 password and localized-key are saved to the configuration file using AES256-CBC encryption.

Example

The following command configures the user guest on the local SNMP Engine with security level noauth (no authentication and no privacy):

```
configure snmpv3 add user guest
```

The following command configures the user authMD5 to use RSA Data Security, Inc. MD5 Message-Digest Algorithm authentication with the password palertyu:

```
configure snmpv3 add user authMD5 authentication md5 palertyu
```

The following command configures the user authShapriv to use SHA authentication with the hex key shown below, the privacy password palertyu, and volatile storage:

```
configure snmpv3 add user authShapriv authentication sha hex  
01:03:04:05:01:05:02:ff:ef:cd:12:99:34:23:ed:ad:ff:ea:cb:11 privacy palertyu volatile
```

History

This command was first available in ExtremeXOS 10.1.

The hex_user_name parameter was added in ExtremeXOS 11.0.

Support for 3DES and AES was added in ExtremeXOS 12.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure snmpv3 add user clone-from

```
configure snmpv3 add user [[hex hex_user_name] | user_name] {engine-
  id engine_id} clone-from [[hex hex_user_name] | user_name] {engine-id
  clone_from_engine_id}
```

Description

Creates a new user by cloning from an existing SNMPv3 user.

Syntax Description

<i>hex_user_name</i>	Specifies the user name to add or to clone from. The value is to be supplies as a colon separated string of hex octets.
<i>user_name</i>	Specifies the user name to add or to clone from in ASCII format.
engine-id	<i>SNMP</i> engine ID
<i>engine_id</i>	Engine ID of the user to be added in hexadecimal format. Default: local engine ID"; type="ostring_t
<i>clone_from_engine_id</i>	Engine ID of the user to be cloned in hexadecimal (Default: local engine ID"); type="ostring_t

Default

N/A.

Usage Guidelines

Use this command to create a new user by cloning an existing one. After you have successfully cloned the new user, you can modify its parameters using the following command:

```
configure snmpv3 add user [[hex hex_user_name] | user_name]
{authentication [md5 | sha] [hexhex_auth_password | auth_password]}
{privacy {des | 3des | aes {128 | 192 | 256}} [[hexhex_priv_password]
| priv_password]} }{volatile}
```

Users cloned from the default users will have the storage type of non-volatile. The default names are: admin, initial, initialmd5, initialsha, initialmd5Priv, initialshaPriv.

Example

The following command creates a user cloneMD5 with same properties as the default user initialmd5. All authorization and privacy keys will initially be the same as with the default user initialmd5.

```
configure snmpv3 add user cloneMD5 clone-from initialmd5
```

The following command adds a remote user named nmsuser2 belonging to the SNMP engine with engine-id 11:22:33 by cloning another remote user named nmsuser1 belonging to the SNMP engine with engine id AA:BB::CC:

```
conf snmpv3 add user nmsuser2 engine-id 11:22:33 clone-from nmsuser1 engine-id AA:BB:CC
```

History

This command was first available in ExtremeXOS 10.1.

The hex_user_name parameter was added in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure snmpv3 delete access

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

Description

Deletes access rights for a group.

Syntax Description

all-non-defaults	Specifies that all non-default (non-permanent) security groups are to be deleted.
<i>hex_group_name</i>	Specifies the group name to be deleted. The value is to be supplies as a colon separated string of hex octets.
<i>group_name</i>	Specifies the group name to be deleted in ASCII format.
sec-model	Specifies the security model to use.
snmpv1	Specifies the SNMPv1 security model.
snmpv2c	Specifies the SNMPv2c security model.
usm	Specifies the SNMPv3 User-based Security Model (USM).
sec-level	Specifies the security level for the group.
noauth	Specifies no authentication (and implies no privacy) for the security level.

authnopriv	Specifies authentication and no privacy for the security level.
priv	Specifies authentication and privacy for the security level.

Default

The default values are:

- sec-model—USM.
- sec-level—noauth.

Usage Guidelines

Use this command to remove access rights for a group. Use the all-non-defaults keyword to delete all the security groups, except for the default groups. The default groups are: admin, initial, vlv2c_ro, vlv2c_rw.

Deleting an access will not implicitly remove the related group to user association from the VACMSecurityToGroupTable. To remove the association, use the following command:

```
configure snmpv3 delete group { [[hex hex_group_name] | group_name] } user
[all-non-defaults | { [[hexhex_user_name] | user_name] {sec-model [snmpv1|
snmpv2c|usm] } } ]
```

Example

The following command deletes all entries with the group name userGroup:

```
configure snmpv3 delete access userGroup
```

The following command deletes the group userGroup with the security model snmpv1 and security level of authentication and no privacy (authnopriv):

```
configure snmpv3 delete access userGroup sec-model snmpv1 sec-level authnopriv
```

History

This command was first available in ExtremeXOS 10.1.

The hex_group_name parameter was added in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure snmpv3 delete community

```
configure snmpv3 delete community [all | { [[hex hex_community_index] |
community_index} | {name [[hex hex_community_name] | community_name}]
```

Description

Deletes an SNMPv3 community entry.

Syntax Description

all	Specifies that all community entries are to be removed.
<i>hex_community_index</i>	Specifies the row index in the snmpCommunityTable. The value is to be supplied as a colon separated string of hex octets.
<i>community_index</i>	Specifies the row index in the snmpCommunityTable in ASCII format.
<i>hex_community_name</i>	Specifies the community name. The value is to be supplied as a colon separated string of hex octets.
<i>community_name</i>	Specifies the community name in ASCII format.

Default

The default entries are public and private.

Usage Guidelines

Use this command to delete an SMMPv3 community in the community MIB.

Example

The following command deletes an entry with the community index comm_index:

```
configure snmpv3 delete community comm_index
```

The following command creates an entry with the community name (hex) of EA:12:CD:CF:AB:11:3C:

```
configure snmpv3 delete community name hex EA:12:CD:CF:AB:11:3C
```

History

This command was first available in ExtremeXOS 10.1.

The hex_community_index and hex_community_name parameters were added in ExtremeXOS 11.0.

The **all-non-defaults** keyword was replaced with the **all** keyword in ExtremeXOS 22.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure snmpv3 delete filter

```
configure snmpv3 delete filter [all | [[hex hex_profile_name] |  
                                profile_name] {subtree object_identifier}]
```

Description

Deletes a filter from a filter profile.

Syntax Description

all	Specifies all filters.
<i>hex_profile_name</i>	Specifies the filter profile of the filter to delete. The value is to be supplied as a colon separated string of hex octets.
<i>profile_name</i>	Specifies the filter profile of the filter to delete in ASCII format.
<i>object_identifier</i>	Specifies the MIB subtree of the filter to delete.

Default

N/A.

Usage Guidelines

Use this command to delete a filter entry from the snmpNotifyFilterTable. Specify all to remove all entries. Specify a profile name to delete all entries for that profile name. Specify a profile name and a subtree to delete just those entries for that filter profile and subtree.

Example

The following command deletes the filters from the filter profile prof1 that reference the MIB subtree 1.3.6.1.4.1:

```
configure snmpv3 delete filter prof1 subtree 1.3.6.1.4.1
```

History

This command was first available in ExtremeXOS 10.1.

The hex_profile_name parameter was added in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure snmpv3 delete filter-profile

```
configure snmpv3 delete filter-profile [all | [hex hex_profile_name |  
    profile_name] {param [hex hex_param_name | param_name}}]
```

Description

Removes the association of a filter profile with a parameter name.

Syntax Description

all	Specifies all filter profiles.
<i>hex_profile_name</i>	Specifies the filter profile name to delete. The value is to be supplied as a colon separated string of hex octets.
<i>profile_name</i>	Specifies the filter profile name to delete in ASCII format.
<i>hex_param_name</i>	Specifies to delete the filter profile with the specified profile name and parameter name. The value is to be supplied as a colon separated string of hex octets.
<i>param_name</i>	Specifies to delete the filter profile with the specified profile name and parameter name in ASCII format.

Default

The default storage type is non-volatile.

Usage Guidelines

Use this command to delete entries from the snmpNotifyFilterProfileTable. This table associates a filter profile with a parameter name. Specify all to remove all entries. Specify a profile name to delete all entries for that profile name. Specify a profile name and a parameter name to delete just those entries for that filter profile and parameter name.

Example

The following example deletes the filter profile prof1 with the parameter name P1:

```
configure snmpv3 delete filter-profile prof1 param P1
```

History

This command was first available in ExtremeXOS 10.1.

The hex_profile_name and hex_param_name parameters were added in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure snmpv3 delete group user

```
configure snmpv3 delete group {[hex hex_group_name] | group_name]} user
[all-non-defaults | {[hex hex_user_name] | user_name} {sec-model
[snmpv1 | snmpv2c | usm] } ]]
```

Description

Deletes a user name (security name) from a group.

Syntax Description

<i>hex_group_name</i>	Specifies the group name to delete or modify. The value is to be supplied as a colon separated string of hex octets.
<i>group_name</i>	Specifies the group name to delete or modify in ASCII format.
all-non-defaults	Specifies that all non-default (non-permanent) users are to be deleted from the group.
<i>hex_user_name</i>	Specifies the user name to delete or modify. The value is to be supplied as a colon separated string of hex octets.
<i>user_name</i>	Specifies the user name to delete or modify in ASCII format.
sec-model	Specifies the security model to use.
snmpv1	Specifies the SNMPv1 security model.
snmpv2c	Specifies the SNMPv2c security model.
usm	Specifies the SNMPv3 User-based Security Model (USM).

Default

The default value for sec-model is USM.

Usage Guidelines

Use this command to remove the associate of a user name with a group.

As per the SNMPv3 RFC, a security name is model independent while a username is model dependent. For simplicity, both are assumed to be same here. User names and security names are handled the same. In other words, if a user is created with the user name username, the security name value is the same, username.

Every group is uniquely identified by a security name and security model. So the same security name can be associated to a group name but with different security models.

The default groups are: admin, initial, v1v2c_ro, v1v2c_rw.

The default users are: admin, initial, initialmd5, initialsha, initialmd5Priv, initialshaPriv.

Example

The following command deletes the user guest from the group UserGroup for the security model snmpv2c:

```
configure snmpv3 delete group UserGroup user guest sec-model snmpv2c
```

The following command deletes the user guest from the group userGroup with the security model USM:

```
configure snmpv3 delete group userGroup user guest
```

History

This command was first available in ExtremeXOS 10.1.

The hex_group_name and the hex_user_name parameters were added in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure snmpv3 delete mib-view

```
configure snmpv3 delete mib-view [all-non-defaults | {[[hex  
    hex_view_name] | view_name] {subtree object_identifier}]]
```

Description

Deletes a MIB view.

Syntax Description

all-non-defaults	Specifies that all non-default (non-permanent) MIB views are to be deleted.
<i>hex_view_name</i>	Specifies the MIB view to delete. The value is to be supplied as a colon separated string of hex octets.
<i>view_name</i>	Specifies the MIB view name to delete in ASCII format.
<i>object_identifier</i>	Specifies a MIB subtree.

Default

N/A.

Usage Guidelines

Use this command to delete a MIB view. Views which are being used by security groups cannot be deleted. Use the all-non-defaults keyword to delete all the MIB views (not being used by security

groups) except for the default views. The default views are: defaultUserView, defaultAdminView, and defaultNotifyView.

Use the `configure snmpv3 add mib-view` command to remove a MIB view from its security group, by specifying a different view.

Example

The following command deletes all views (only the permanent views will not be deleted):

```
configure snmpv3 delete mib-view all-non-defaults
```

The following command deletes all subtrees with the view name AdminView:

```
configure snmpv3 delete mib-view AdminView
```

The following command deletes the view AdminView with subtree 1.3.6.1.2.1.2

```
configure snmpv3 delete mib-view AdminView subtree 1.3.6.1.2.1.2
```

History

This command was first available in ExtremeXOS 10.1.

The `hex_view_name` parameter was added in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure snmpv3 delete notify

```
configure snmpv3 delete notify [{{[[hex hex_notify_name] | notify_name}}
| all-non-defaults]
```

Description

Deletes an entry from the snmpNotifyTable.

Syntax Description

<i>hex_notify_name</i>	Specifies the notify name to add. The value is to be supplied as a colon separated string of hex octets.
<i>notify_name</i>	Specifies the notify name to add in ASCII format.
all-non-defaults	Specifies that all non-default (non-permanent) notifications are to be deleted.

Default

N/A.

Usage Guidelines

Use this command to delete an entry from the snmpNotifyTable. When a notification is to be sent, this table is examined. For the target addresses that have been associated with the tags present in the table, notifications will be sent, based on the filters also associated with the target addresses.

Example

The following command removes the N1 entry from the table:

```
configure snmpv3 delete notify N1
```

History

This command was first available in ExtremeXOS 10.1.

The hex_notify_name parameter was added in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure snmpv3 delete target-addr

```
configure snmpv3 delete target-addr {[[hex hex_addr_name] | addr_name}]  
| all]
```

Description

Deletes SNMPv3 target addresses.

Syntax Description

<i>hex_addr_name</i>	Specifies an identifier for the target address. The value is to be supplied as a colon separated string of hex octets.
<i>addr_name</i>	Specifies a string identifier for the target address.
all	Specifies all target addresses.

Default

N/A.

Usage Guidelines

Use this command to delete an entry in the SNMPv3 snmpTargetAddressTable.

Example

The following command deletes target address named A1:

```
configure snmpv3 delete target-addr A1
```

History

This command was first available in ExtremeXOS 10.1.

The hex_addr_name parameter was added in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure snmpv3 delete target-params

```
configure snmpv3 delete target-params [{[[hex hex_param_name] |  
param_name]} | all]
```

Description

Deletes SNMPv3 target parameters.

Syntax Description

<i>hex_param_name</i>	Specifies the parameter name associated with the target. The value is to be supplied as a colon separated string of hex octets.
<i>param_name</i>	Specifies the parameter name associated with the target in ASCII format.

Default

N/A.

Usage Guidelines

Use this command to delete an entry in the SNMPv3 snmpTargetParamsTable. This table specifies the message processing model, security level, security model, and the storage parameters for messages to any target addresses associated with a particular parameter name.

Example

The following command deletes a target parameters entry named P1:

```
configure snmpv3 delete target-params P1
```

History

This command was first available in ExtremeXOS 10.1.

The hex_param_name parameter was added in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure snmpv3 delete user

```
configure snmpv3 delete user [all | [[hex hex_user_name] | user_name]
                               {engine-id engine_id}]
```

Description

Deletes an existing SNMPv3 user.

Syntax Description

all	Specifies that all users are to be deleted.
<i>hex_user_name</i>	Specifies the user name to delete. The value is to be supplied as a colon separated string of hex octets.
<i>user_name</i>	Specifies the user name to delete.
engine-id	<i>SNMP</i> engine ID
<i>engine-id</i>	Engine ID in hexadecimal (Default: local engine ID); type="ostring_t"

Default

N/A.

Usage Guidelines

Use this command to delete an existing user.

Deleting users does not implicitly remove the related group-to-user association from the VACMSecurityToGroupTable. To remove the association, use the following command:

```
configure snmpv3 delete group {[[hex hex_group_name] | group_name]}
user [all-non-defaults | {[[hex hex_user_name] | user_name] {sec-model
[snmpv1|snmpv2c|usm]}}]
```

Example

The following command deletes all users:

```
configure snmpv3 delete user all
```

The following command deletes the user "guest":

```
configure snmpv3 delete user guest
```

The following command deletes a remote user named "ambiguoususer" with engine id 11:22:33:

```
configure snmpv3 delete user ambiguoususer engine-id 11:22:33
```

History

- This command was first available in ExtremeXOS 10.1.
- The **hex_user_name** parameter was added in ExtremeXOS 11.0.
- The **engine_id** keyword was added in ExtremeXOS 15.4.
- The **all-non-default** keyword was replaced with the **all** keyword in ExtremeXOS 22.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure snmpv3 engine-boots

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

Description

Configures the SNMPv3 Engine Boots value.

Syntax Description

(1-2147483647)	Specifies the value of engine boots.
----------------	--------------------------------------

Default

N/A.

Usage Guidelines

Use this command if the Engine Boots value needs to be explicitly configured. Engine Boots and Engine Time will be reset to one (1) if the Engine ID is changed. Engine Boots can be set to any desired value, but will latch on its maximum, 2147483647.

Example

The following command configures Engine Boots to 4096:

```
configure snmpv3 engine-boots 4096
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure snmpv3 engine-id

```
configure snmpv3 engine-id hex_engine_id
```

Description

Configures the SNMPv3 snmpEngineID.

Syntax Description

<i>hex_engine_id</i>	Specifies the colon delimited hex octet that serves as part of the snmpEngineID (5-32 octets).
----------------------	--

Default

The default snmpEngineID is the device MAC address.

Usage Guidelines

Use this command if the snmpEngineID needs to be explicitly configured. The first four octets of the ID are fixed to 80:00:07:7C, which represents Extreme Networks Vendor ID. Once the snmpEngineID is changed, default users are reverted back to their original passwords/keys, while non-default users are removed from the device.

Example

The following command configures the snmpEngineID to be 80:00:07:7C:00:0a:1c:3e:11:

```
configure snmpv3 engine-id 00:0a:1c:3e:11
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure snmpv3 target-addr retry

```
configure snmpv3 target-addr [[hex hex_addr_name] | addr_name] retry
                               retry_count
```

Description

Configures SNMPv3 INFORM notification retries.

Syntax Description

<i>hex_addr_name</i>	Specifies a address name in hexadecimal format.
<i>addr_name</i>	Specifies the address name in ASCII format.
<i>retry_count</i>	Specifies the maximum number of times to resend an SNMPv3 inform.

Default

The retry default is 3.

Usage Guidelines

Use this command to configure the number of times an SNMPv3 INFORM message is to be resent to the (notification responder) manager when a response has not been received.

Example

The following command configures a retry count of 5 for the target address A1:

```
configure snmpv3 target-addr A1 retry 5
```

History

This command was first available in ExtremeXOS 12.5.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure snmpv3 target-addr timeout

```
configure snmpv3 target-addr [[hex hex_addr_name] | addr_name] timeout
                               timeout_val
```


Description

Configures the SNMPv3 INFORM notification timeout.

Syntax Description

<i>hex_addr_name</i>	Specifies the address name in hexadecimal format.
<i>addr_name</i>	Specifies the address name in ASCII format.
<i>timeout_val</i>	Specifies the number of seconds.

Default

The timeout value default is 15 seconds.

Usage Guidelines

Use this command to configure how many seconds to wait for a response before resending an SNMPv3 INFORM.

Example

The following command configures a timeout value of 20 seconds for the target address A1:

```
configure snmpv3 target-addr A1 timeout 20
```

History

This command was first available in ExtremeXOS 12.5.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure sntp-client

```
configure sntp-client [primary | secondary] host-name-or-ip {vr vr_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.
<i>host-name-or-ip</i>	Specifies a host name or IPv4 address or IPv6 address.

vr	Specifies use of a virtual router.NOTE: User-created VRs are supported only on the platforms listed for this feature in the Switch Engine 32.3 Feature License Requirements document.
<i>vr_name</i>	Specifies the name of a virtual router.

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 example configures a primary NTP server:

```
configure sntp-client primary 10.1.2.2
```

The following example configures the primary NTP server to use the management virtual router [VR-Mgmt](#):

```
configure sntp-client primary 10.1.2.2 vr VR-Mgmt
```

History

This command was first available in ExtremeXOS 10.1.

The **vr** *vr_name* option was added in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5420, 5520, 5720 series switches.

configure sntp-client update-interval

```
configure sntp-client update-interval update-interval
```

Description

Configures the interval between polls for time information from [SNTP](#) servers.

Syntax Description

<i>update-interval</i>	Specifies an interval in seconds.
------------------------	-----------------------------------

Default

64 seconds.

Usage Guidelines

None.

Example

The following command configures the interval timer:

```
configure sntp-client update-interval 30
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ssh2 access-profile

```
configure ssh2 access-profile [ access_profile | [[add rule] [first |  
[before | after] previous_rule]] | delete rule | none]
```

Description

Configures SSH2 to use an [ACL](#) policy or ACL rule for access control.

Syntax Description

<i>access_profile</i>	Specifies an ACL policy.
add	Specifies that an ACL rule is to be added to the SSH2 port.
<i>rule</i>	Specifies an ACL rule.
first	Specifies that the new rule is to be added before all other rules.
before	Specifies that the new rule is to be added before a previous rule.
after	Specifies that the new rule is to be added after a previous rule.
<i>previous_rule</i>	Specifies an existing rule in the application.
delete	Specifies that one particular rule is to be deleted.
none	Specifies that all the rules or a policy file is to be deleted.

Default

N/A.

Usage Guidelines

You must be logged in as administrator to configure SSH2 parameters.

- Implement an ACL policy file that permits or denies a specific list of IP addresses and subnet masks for the SSH2 port. You must create the ACL policy file before you can use this command. If the ACL policy file does not exist on the switch, the switch returns an error message indicating that the file does not exist.

In the ACL policy file for SSH2, the “source-address” field is the only supported match condition. Any other match conditions are ignored.

Use the none option to remove a previously configured ACL.

Policy files can also be configured using the `enable ssh2` command.

- Add an ACL rule to the SSH2 application through this command. Once an ACL is associated with SSH2, all the packets that reach an SSH2 module are evaluated with this ACL and appropriate action (permit or deny) is taken, as is done using policy files.

The permit or deny counters are also updated accordingly regardless of whether the ACL is configured to add counters. To display counter statistics, use the `show access-list counters process` command.

Only the following match conditions and actions are copied to the client memory. Others that may be in the rule are not copied.

Match conditions:

- Source-address—IPv4 and IPv6
- Actions—Permit or Deny

When adding a new rule, use the first, before, and after previous_rule parameters to position it within the existing rules.

If the SSH2 traffic does not match any of the rules, the default behavior is deny. To permit SSH2 traffic that does not match any of the rules, add a permit all rule at the end of the rule list.

Creating an ACL Policy File

To create an ACL policy file, use the `edit policy` command. For more information about creating and implementing ACL policy files, see [Policy Manager](#) and [ACLs](#) in the *Switch Engine 32.3 User Guide*.

If you attempt to implement a policy that does not exist on the switch, an error message similar to the following appears:

```
Error: Policy /config/MyAccessProfile.pol does not exist on file system
```

If this occurs, make sure the policy you want to implement exists on the switch. To confirm the policies on the switch, use the `ls` command. If the policy does not exist, create the ACL policy file.

Example

The following example applies the ACL MyAccessProfile_2 to SSH2:

```
configure ssh2 access-profile MyAccessProfile_2
```

The following example copies the ACL rule, DenyAccess to the SSH2 application in first place:

```
configure ssh2 access-profile add DenyAccess first
```

The following example removes the association of a single rule from the SSH2 application:

```
configure ssh2 access-profile delete DenyAccess
```

The following example removes the association of all ACL policies and rules from the SSH2 application:

```
configure ssh2 access-profile none
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ssh2 dh-group

```
configure ssh2 dh-group minimum [1 | 14 | 16 | 18]
```

Description

Configures the minimal supported Diffie-Hellman group.

Syntax Description

dh-group	Configures the Diffie-Hellman group. Used for cryptographic key exchange. Higher groups are stronger.
minimum	Configures minimal supported Diffie-Hellman group to avoid using weaker groups.
1	Supports Diffie-Hellman group 1 (1,024 bit), 14 (2,048 bit), 16 (4,096 bit), and 18 (8,192 bit).
14	Supports group 14 (2,048 bit), 16 (4,096 bit), and 18 (8,192 bit). Default.
16	Supports Diffie-Hellman group 16 (4,096 bits) and 18 (8,192 bits).
18	Supports only Diffie-Hellman group 18 (8,192 bits).

Default

The minimal supported Diffie-Hellman group is 14. This means that Diffie-Hellman groups 14, 16, and 18 are supported by default.

Usage Guidelines

Openssh-7.5p1 supports Diffie-Hellman group 1, 14, 16, and 18 as part of the key exchange algorithms. By default, Diffie-Hellman group 14, 16, and 18 are supported.

To revert back to using Diffie-Hellman group 1 (in addition to Diffie-Hellman group 14, 16, and 18), set the minimal support group to Diffie-Hellman group1.

The server picks the first entry from the client proposal and matches it with its own proposal. If there is no match, the server picks the next entry from the client proposal and so on. If no match is found, the connection is rejected.

Example

The following example configures Diffie-Hellman group 16 as the minimum supported Diffie-Hellman group.

```
configure ssh2 dh-group minimum 16
```

History

This command was first available in ExtremeXOS 22.1.

Support for Diffie-Hellman groups 16 and 18 was added in ExtremeXOS 22.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ssh2 disable cipher mac

```
configure ssh2 disable [cipher [cipher | all] | mac [ mac | all ]]
```

Description

Disables ciphers/Message Authentication Codes (MACs) for use with SSHv2.

Syntax Description

cipher	Specifies cipher to disable for the encrypting session.
<i>cipher</i>	Specific cipher name to disable.
all	Specifies all ciphers/MACs available in current mode.
mac	Specifies MACs to disable for the encrypting session.
<i>mac</i>	Specific MAC name to disable.

Default

None.

Example

The following example disables cipher "aes256-ctr":

```
configure ssh2 disable cipher "aes256-ctr"
```

History

This command was first available in ExtremeXOS 22.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ssh2 disable pk-alg

```
configure ssh2 disable {pk-alg [pkalg_name | all]}
```

Description

Disables DSA/RSA X509v3 public key algorithms.

Syntax Description

pk-alg	Specifies disabling DSA/RSA X509v3 public key algorithms.
<i>pkalg_name</i>	Specifies which algorithm to disable: "ssh-dss" "ssh-rsa" "x509v3-sign-dss" "x509v3-sign-rsa"
all	Specifies disabling all public key algorithms available.

Default

By default all the algorithms are enabled.

Example

The following example disables the ssh-dss algorithm:

```
configure ssh2 disable pk-alg ssh-dss
```

History

This command was first available in ExtremeXOS 22.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ssh2 enable cipher mac

```
configure ssh2 enable [cipher [cipher | all] | mac [mac | all]]
```

Description

Configures the required ciphers/Message Authentication Codes (MACs) with SSHv2.

Syntax Description

cipher	Specifies cipher to use for encrypting the session.
<i>cipher</i>	Cipher name for encrypting session.
all	Specifies all ciphers/MACs available in current mode.
mac	Specifies MACs to use for encrypting the session.
<i>mac</i>	MAC name for encrypting session.

Default

In Default mode, the following ciphers/MACs are *disabled* by default:

- **Ciphers:** 3des-cbc, aes128-cbc, aes192-cbc, aes256-cbc, rijndael-cbc@lysator.liu.se
- **MACs:** hmac-md5, hmac-md5-96, hmac-md5-etm@openssh.com, hmac-md5-96-etm@openssh.com, hmac-sha1-96, hmac-sha1-96-etm@openssh.com

In Default mode, the following ciphers/MACs are *enabled* by default:

- **Ciphers:** aes128-ctr, aes192-ctr, aes256-ctr, chacha20-poly1305@openssh.com
- **MACs:** hmac-sha1-etm@openssh.com, hmac-sha2-256-etm@openssh.com, hmac-sha2-512-etm@openssh.com, hmac-sha1, hmac-sha2-256, hmac-sha2-512.



Note

The following ciphers and MAC are no longer supported: arcfour, arcfour128, arcfour256, blowfish-cbc, cast128-cbc, hmac-ripemd160.

Example

The following example enables cipher "aes256-ctr" for the encrypting the session:

```
# configure ssh2 enable cipher "aes256-ctr"
```

History

This command was first available in ExtremeXOS 22.1.

Unsupported ciphers/macs removed due to SSH2 upgrade in ExtremeXOS 30.7.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ssh2 enable pk-alg

```
configure ssh2 enable {pk-alg [pkalg_name | all]}
```

Description

Enables DSA/RSA X509v3 public key algorithms.

Syntax Description

pk-alg	Specifies enabling DSA/RSA X509v3 public key algorithms.
<i>pkalg_name</i>	Specifies which algorithm to enable: "ssh-dss" "ssh-rsa" "x509v3-sign-dss" "x509v3-sign-rsa"
all	Specifies enabling all public key algorithms available.

Default

ssh-dss is *disabled* by default.

ssh-rsa, x509v3-sign-rsa, x509v3-sign-dss are *enabled* by default.

Usage Guidelines

This public key algorithm configuration is used for the user key only—not for the host key. For a user key, ssh-dss algorithm is supported, but disabled by default. However, for host key, ssh-dss algorithm is not supported for both server and client. For backward compatibility it is supported in the server only during a switch image upgrade if this algorithm is present in earlier release.

Example

The following example enables the ssh-dss algorithm:

```
configure ssh2 enables pk-alg ssh-dss
```

History

This command was first available in ExtremeXOS 22.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ssh2 idletimeout

```
configure ssh2 idletimeout [none | minutes]
```

Description

This command configures idle-timeout for SSH/SFTP connections.

Syntax Description

none	Idle timeout disabled.
<i>minutes</i>	Timeout value in minutes. Range is 1 to 240.

Default

60 minutes.

Usage Guidelines

If you enable the idle timer using the enable idletimeout command, the SSH2 connection times out after 20 minutes of inactivity by default. If you disable the idle timer using the disable idletimeout command, the SSH2 connection times out after 60 minutes of inactivity by default. This timeout value can be modified using the command “configure ssh2 idletimeout <minutes> wherein <minutes> can be from 1 to 240 “. This ssh idle timer is applicable for SFTP connections as well.

Example

Configured ssh idle timeout is displayed in “show management” output:

```
# show management
CLI idle timeout           : Enabled (2 minutes)
CLI max number of login attempts :      3
CLI max number of sessions   :      8
CLI paging                  : Enabled (this session only)
CLI space-completion         : Disabled (this session only)
CLI configuration logging     : Disabled
CLI journal size             : 100
CLI password prompting only   : Disabled
CLI scripting                 : Disabled (this session only)
CLI scripting error mode      : Ignore-Error (this session only)
CLI persistent mode           : Persistent (this session only)
CLI prompting                 : Enabled (this session only)
CLI screen size               : 32 Lines 112 Columns (this session only)
CLI refresh                   : Enabled
Telnet access                 : Enabled (tcp port 23 vr all)
                             : Access Profile : not set
SSH access                    : Enabled (Key valid, tcp port 22 vr all)
                             : Access Profile : not set
SSH2 idle timeout             : 20 minutes
Web access                    : Enabled (tcp port 80)
                             : Access Profile : not set
Total Read Only Communities   : 1
Total Read Write Communities  : 1
RMON                          : Disabled
SNMP access                   : Enabled
                             : Access Profile : not set
SNMP Notifications            : Enabled
SNMP Notification Receivers    : None
SNMP stats:      InPkts 0      OutPkts 0      Errors 0      AuthErrors 0
                  Gets  0      GetNexts 0      Sets  0      Drops  0
```

```
SNMP traps:      Sent    0      AuthTraps Enabled
SNMP inform:     Sent    0      Retries    0      Failed 0
```

History

This command was first available in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ssh2 key

```
configure ssh2 key {pregenerated}
```

Description

Generates the Secure Shell 2 (SSH2) host key. This command is used to regenerate a host key, if there is already one existing.

Syntax Description

pregenerated	indicates that the SSH2 host key is already available with the user.
---------------------	--

Default

The switch generates a key for each SSH2 session.

Usage Guidelines

Secure Shell 2 (SSH2) is a feature of ExtremeXOS that allows you to encrypt 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. Configuration, policy, image, and public key files may also be transferred to the switch using the Secure Copy Program (SCP2).

To enable SSH2, use the `enable ssh2` command.

A host key must be generated before the switch can accept incoming ssh connections. This can be done by the switch using the commands "enable ssh2" (if ssh is not enabled previously) or "configure ssh2 key pregenerated" (if you wish to use a pregenerated key as the host key).

If you elect to have the key generated, the key generation process can take up to one minute, and cannot be canceled after it has started. For the switch to use the newly generated key the `exsshd` process needs to be restarted using the command `restart process [class cname | name {msm slot}]` with "exsshd" as the *name*.

To use a key that has been previously created, use the **pregenerated** keyword. Use the `show ssh2 private-key` command to list and copy the previously generated key. Then use the `configure ssh2 key {pregenerated}` command where “pregenerated” represents the key that you paste.

**Note**

In ExtremeXOS 22.5 and later, ssh-dss (DSA) host key is not supported in both server and client. For backward compatibility, it is supported in server only during a switch image upgrade if this algorithm is present in earlier release.

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.

To view the status of SSH2 on the switch, use the `show management` command. The `show management` command displays information about the switch including the enable/disable state for SSH2 sessions, whether a valid key is present, and the TCP port and virtual router that is being used.

Example

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

```
configure 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 begins the process.

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

```
configure ssh2 key pregenerated <pre-generated key>
```

Enter the previously-generated key (you can copy and paste it from the saved configuration file; a part of the key pattern is similar to 2d:2d:2d:2d:20:42:45:47:).

History

This command was first available in the ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ssh2 login-grace-timeout

```
configure ssh2 login-grace-timeout seconds
```

Description

For the SSH server, configures a timeout period for a login attempt.

Syntax Description

ssh2	Designates changing SSHv2 configuration.
login-grace-timeout	Designates having the SSHv2 server disconnects after this time if the user has not completed login attempt.
<i>seconds</i>	Sets the time in seconds for the disconnect timeout period. The default is 120 seconds, and the range is 60–120 seconds.

Default

By default, the timeout period is 120 seconds.

Usage Guidelines

To view the current timeout period setting , use the command `show ssh2`.

Example

The following example sets the timeout period ot 100 seconds:

```
# configure ssh2 login-grace-timeout 100
```

History

This command was first available in ExtremeXOS 30.7.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ssh2 rekey

```
configure ssh2 rekey [time-interval [time_interval | none] | data-limit [data_size | default]]
```

Description

Sets SSHv2 session rekeying interval by specifying a time interval value and/or amount of transferred data.

Syntax Description

ssh2	Specifies setting SSHv2 behavior.
rekey	Specifies rekey request interval for SSH connection.
time-interval	Sets rekey time interval.
<i>time_interval</i>	Specifies rekey time interval value in minutes. Valid range 1 to 1,440.
none	Specifies no time limit for rekey interval (default).
data-limit	Specifies rekey interval in terms of amount of data transferred.
<i>data_size</i>	Sets data transfer limit in MB. Valid range is 1 to 4,096 MB.
default	Sets the data limit to the default specified by the cipher. Values range between 1GB and 4GB. This is the default setting.

Default

If nothing is specified, the rekey time interval is set to none, and the data limit is specified by the cipher in use.

Usage Guidelines

You can set both a time limit and a data limit for the rekey interval. Your selections for rekeying appear in the output of the `show ssh2` command.

Example

The following example sets the SSHv2 rekey time interval to one hour (60 mins):

```
configure ssh2 rekey time-interval 60
```

History

This command was first available in ExtremeXOS 22.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ssh2 secure-mode

```
configure ssh2 secure-mode [on | off]
```

Description

This command (secure-mode on) disables the weak ciphers and macs in SSH server and client.

Syntax Description

on	Enable all supported algorithms.
off	Enable only compliance algorithms.

Default

Off.

Usage Guidelines

After enabling secure-mode:

- For communication, SSH server uses a new secure-mode list made each for ciphers and macs.
- For SSH client, EPM is notified to change the bit dedicated to SSH secure-mode, which hides the weak ciphers and macs from SSH client CLI commands.

Example

```
configure ssh2 secure-mode on

show management
CLI idle timeout           : Disabled
CLI max number of login attempts : 3
CLI max number of sessions  : 8
CLI paging                 : Enabled (this session only)
CLI space-completion       : Disabled (this session only)
CLI configuration logging   : Enabled
CLI password prompting only : Disabled
CLI RADIUS cmd authorize tokens : 2
CLI scripting              : Disabled (this session only)
CLI scripting error mode    : Ignore-Error (this session only)
CLI persistent mode        : Persistent (this session only)
CLI prompting              : Enabled (this session only)
CLI screen size            : 24 Lines 80 Columns (this session only)
CLI refresh                : Enabled
Telnet access              : Enabled (tcp port 23 vr all)
                           : Access Profile : not set
SSH access                 : Enabled (Key valid, tcp port 22 vr all)
                           : Secure-Mode   : On
                           : Access Profile : not set
SSH2 idle time             : 60 minutes
Web access                 : Enabled (tcp port 80)
                           : Access Profile : not set
Total Read Only Communities : 1
Total Read Write Communities : 1
RMON                      : Disabled
SNMP access               : Enabled
                           : Access Profile : not set
SNMP Notifications        : Enabled
SNMP Notification Receivers : None
SNMP stats:               InPkts 0      OutPkts 0      Errors 0      AuthErrors
0
                           Gets 0      GetNexts 0      Sets 0      Drops 0
SNMP traps:               Sent 0      AuthTraps Enabled
SNMP inform:              Sent 0      Retries 0      Failed 0
```

History

This command was first available in ExtremeXOS 21.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ssh2 x509v3 ocs

```
configure ssh2 x509v3 ocs [on | off]
```

Description

Enables or disables Online Certificate Status Protocol (OCSP) check for SSH2 x509v3 authentication.

Syntax Description

x509v3	Specifies x509v3 certificate-based authentication.
ocs	Specifies configuring OCSP for real-time certificate revocation status checking.
on	Enables OCSP (default).
off	Disables OCSP.

Default

By default, OCSP is enabled.

Usage Guidelines

While you can disable OCSP, it is not recommended because no certificate revocation status check is performed.

Example

The following example enables OCSP check for SSH2 x509v3 servers.

```
# configure ssh2 x509v3 ocs on
```

History

This command was first available in ExtremeXOS 32.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ssh2 x509v3 ocsf nonce

```
configure ssh2 x509v3 ocsf nonce [on | off]
```

Description

Enables or disables the Online Certificate Status Protocol (OCSP) nonce for SSH2 x509v3 authentication.

Syntax Description

x509v3	Specifies x509v3 certificate-based authentication.
ocsf	Specifies configuring OCSP for real-time certificate revocation status checking.
nonce	Specifies to cryptographically bind an OCSP request and an OCSP response with the extension <code>id-pkix-ocsf-nonce</code> to prevent replay attacks.
on	Specifies to include the <code>id-pkix-ocsf-nonce</code> extension in the OCSP request and response.
off	Specifies to exclude the extension (default).

Default

Off.

Usage Guidelines

Example

The following example configures nonce:

```
# configure ssh2 x509v3 ocsf nonce on
```

History

This command was first available in ExtremeXOS 32.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ssh2 x509v3 ocsf override

```
configure ssh2 x509v3 ocsf override [url | none]
```

Description

This command configures one HTTP Online Certificate Status Protocol (OCSP) override URL for an SSH2 x509v3 authentication.

Syntax Description

x509v3	Specifies x509v3 certificate-based authentication.
ocsp	Specifies the OCSP attribute.
override	Specifies to override the OCSP server in the AuthorityInformationAccess section of a syslog server's certificate.
<i>url</i>	Specifies the URL of the OCSP override server. Default port is 80.
none	Specifies to remove the OCSP override URL configuration (default).

Default

None.

Usage Guidelines

Only HTTP is supported with either FQDN or IP.

Example

The following example configures an override URL of `http://sshocsp:2023`:

```
# configure ssh2 x509v3 ocsp override http://sshocsp:2023
```

History

This command was first available in ExtremeXOS 32.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ssh2 x509v3 ocsp signer

```
configure ssh2 x509v3 ocsp signer ocsp-nocheck [on | off]
```

Description

Enables or disables Online Certificate Status Protocol (OCSP) signer's ocsp-nocheck for SSH2 x509v3 authentication.

Syntax Description

x509v3	Specifies x509v3 certificate-based authentication.
ocsp	Specifies configuring OCSP for real-time certificate revocation status checking.
ocsp-nocheck	Specifies the extension <code>id-pkix-ocsp-nocheck</code> . If present in the OCSP signer's certificate, then it is trusted for its lifetime.
on	Specifies to override the <code>id-pkix-ocsp-nocheck</code> extension in the OCSP signer's certificate and forces the extension as if it is present.
off	Specifies to behave per the extension's presence in the OCSP signer's certificate. If not present and the OCSP signer is not root CA, then the whole OCSP will fail (default).
signer	Specifies the OCSP signer that signs the OCSP response.

Default

Off.

Usage Guidelines

Example

The following example enables OCSP signer's `ocsp-nocheck` for a SSH2 x509v3 server.

```
# configure ssh2 x509v3 ocsp signer ocsp-nocheck on
```

History

This command was first available in ExtremeXOS 32.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ssh2 x509v3 radius-password-auth

```
configure ssh2 x509v3 radius-password-auth [on | off]
```

Description

Enables or disables password authentication using RADIUS for SSH2 x509v3 publication-key authentication.

Syntax Description

x509v3	Specifies x509v3 public-key authentication.
radius-password-auth	Specifies to obtain the password from the user and authenticate it using RADIUS server for X509v3 public-key authentication.
on	Specifies to enable password authentication using RADIUS for X509v3 public-key authentication.
off	Specifies to disable password authentication using RADIUS for X509v3 public-key authentication. (Default).

Default

Off.

Usage Guidelines

Example

The following example enables password authentication using RADIUS for X509v3 public-key authentication.

```
# configure ssh2 x509v3 radius-password-auth on
Note: When turned on, user provides password for RADIUS authentication.

If RADIUS is not configured, local authentication is used.
```

History

This command was first available in ExtremeXOS 32.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ssh2 x509v3 username overwrite

```
configure ssh2 x509v3 username overwrite [on | off]
```

Description

Enables or disables authentication username configuration to use the Principal Name in the certificate as the username.

Syntax Description

x509v3	Specifies x509v3 public-key authentication.
username	Specifies X509v3 authentication username configuration.

overwrite	Specifies that when radius-password-auth is turned on, to use the Principal Name in the certificate as the username.
on	Specifies to enable X509v3 authentication username configuration.
off	Specifies to disable X509v3 authentication username configuration. (Default).

Default

Off.

Usage Guidelines

Example

The following example enable using the Principal Name in the certificate as the username.

```
# configure ssh2 x509v3 username overwrite on
```

Note: This command is applicable only if X509v3 'radius-password-auth' command is turned on.

History

This command was first available in ExtremeXOS 32.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ssh2 x509v3 username strip-domain

```
configure ssh2 x509v3 username strip-domain [on | off]
```

Description

Enables or disables authentication username configuration to strip the domain name for SSH2 x509v3 public-key authentication.

Syntax Description

x509v3	Specifies x509v3 public-key authentication.
username	Specifies X509v3 authentication username configuration.
strip-domain	Specifies that when radius-password-auth and username overwrite are turned on, to strip the domain name from the username.

on	Specifies to enable X509v3 authentication username configuration.
off	Specifies to disable X509v3 authentication username configuration. (Default).

Default

Off.

Usage Guidelines

Example

The following example enables X509v3 authentication username configuration to strip the domain name from the username:

```
# configure ssh2 x509v3 username strip-domain on
```

Note: This command is applicable only if X509v3 'radius-password-auth' and username 'overwrite' commands are turned on.

History

This command was first available in ExtremeXOS 32.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ssh2 x509v3 username use-domain

```
configure ssh2 x509v3 username use-domain [domain_name | none]
```

Description

Enables or disables authentication username configuration with a domain name using RADIUS for SSH2 x509v3 public-key authentication.

Syntax Description

x509v3	Specifies x509v3 public-key authentication.
username	Specifies X509v3 authentication username configuration.
use-domain	Specifies that when radius-password-auth , username overwrite , and strip-domain are turned on, to use the configured domain name as the username.
<i>domain_name</i>	Specifies the domain name to be added to the username.
none	Specifies to remove the use-domain configuration. (Default).

Default

None.

Usage Guidelines

Example

The following example enable authentication username configuration with a domain name of 'abcdef.com':

```
# configure sshd2 x509v3 username use-domain abcdef.com
```

Note: This command is applicable only if X509v3 'radius-password-auth', username 'overwrite' and 'strip-domain' commands are turned on.

History

This command was first available in ExtremeXOS 32.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure sshd2 user-key add user

```
configure sshd2 user-key key_name add user user_name
```

Description

Associates a user to a key.

Syntax Description

<i>key_name</i>	Specifies the name of the public key.
<i>user_name</i>	Specifies the name of the user.

Default

N/A.

Usage Guidelines

This command associates (or binds) a user to a key. Pressing **TAB** at the end of the command lists existing account names.

Example

The following example binds the key `id_dsa_2048` to user `admin`:

```
configure sshd2 user-key id_dsa_2048 add user admin
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure sshd2 user-key delete user

```
configure sshd2 user-key key_name delete user user_name
```

Description

Disassociates a user to a key.

Syntax Description

<i>key_name</i>	Specifies the name of the public key.
<i>user_name</i>	Specifies the name of the user.

Default

N/A.

Usage Guidelines

This command disassociates (or unbinds) a user to a key. Pressing **TAB** at the end of the command shows a list of users attached to the key.

Example

The following example unbinds the key `id_dsa_2048` from user `admin`:

```
configure sshd2 user-key id_dsa_2048 delete user admin
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ssl certificate hash-algorithm

```
configure ssl certificate hash-algorithm hash_algorithm
```

Description

This command configures the hash algorithm.

Syntax Description

ssl	SSL.
certificate	Certificate.
hash-algorithm	Hash algorithm to use (Default SHA-512).
<i>hash_algorithm</i>	Name of hash algorithm to use (Default SHA-512).

Default

SHA-512 algorithm.

Usage Guidelines

Use this command to configure the hash algorithm. Once configured, this configured algorithm will be used for the next certificate creation. Previously *MD5* was the only hashing algorithm available. As of ExtremeXOS 16.1, the default has been changed to more secure SHA-512 algorithm. If you prefer the older version, you can configure to the least secure MD5 hashing algorithm.

Example

The following example displays the show ssl output with the SHA-512 algorithm configured:

```
X460G2-48t-10G4.5 # show ssl
HTTPS Port Number: 443 (Enabled)
Signature Algorithm configured: SHA-512 with RSA Encryption
Private Key matches the Certificate's public key.
RSA Key Length: 1024
Certificate:
  Data:
    Version: 3 (0x2)
    Serial Number: 0 (0x0)
    Signature Algorithm: md5WithRSAEncryption
    Issuer: C=IN, O=ext, CN=ext
    Validity
      Not Before: Dec  7 21:52:53 2014 GMT
      Not After : Dec  7 21:52:53 2015 GMT
    Subject: C=IN, O=ext, CN=ext
```

History

This command was first available in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ssl certificate pregenerated

```
configure ssl certificate pregenerated { {csr-cert}pregenerated {ocsp {on
| off}}}
```

Description

Obtains the pre-generated certificate from the user.

Syntax Description

ssl	SSL.
certificate	Certificate.
csr-cert	Specifies the SSL/TLS certificate signed through CSR generated by switch. Trust chain verification performed during configuration. Only use this option for CSR-signed certificates.
pregenerated	Specifies already having a certificate or private key in Privacy Enhanced Mail (PEM) format.
ocsp	Specifies Online Certificate Status Protocol (OCSP). This option is only available if you have selected CSR-signed certificates.
on	Enables OCSP for SSL/TLS certificate signed through CSR generated by the switch.
off	Disables OCSP for SSL/TLS certificate signed through CSR generated by the switch (default).

Default

For CSR-signed certificates, OCSP is off by default.

Usage Guidelines

You must upload or generate a certificate for SSL server use. With this command, you copy and paste the certificate into the command line followed by a blank line to end the command. The following security algorithms are supported:

- RSA for public key cryptography (generation of certificate and public-private key pair, certificate signing). RSA key size between 2,048 and 4,096 bits.
- Symmetric ciphers (for data encryption): RC4, DES, and 3DES.
- Message Authentication Code (MAC) algorithms: RSA Data Security, Inc. [*MD5*](#) Message-Digest Algorithm and SHA.

This command is also used when downloading or uploading the configuration. Do not modify the certificate stored in the uploaded configuration file because the certificate is signed using the issuer's private key.

The certificate and private key file should be in PEM format and generated using RSA as the cryptography algorithm.

Only use the **csr-cert** option for CSR-signed certificates.

When a certificate is imported using this **csr-cert** option, mandatory trust chain verification and optional revocation check is performed. For a successful import, both verifications should pass. ExtremeXOS supports the revocation checking using the OCSP library. During the import of the switch certificate, if it is with **csr-cert** option, then if the trust chain verification passes, then the revocation status of the switch certificate and a maximum of 5 intermediate CA certificates (total of 6 certificates). When OCSP on is chosen, a revocation check is performed. The certificate is accepted only when revocation status is good for all certificates (switch and a maximum of 5 intermediate CA). If the revocation status is anything other than good (including unable to connect, no response, revoked, unknown) for any of the above certificates, then that certificate import is rejected. It can be imported though, by selecting OCSP as off.

Example

The following command obtains the pre-generated certificate from the user:

```
configure ssl certificate pregenerated
```

Next, you open the certificate, and then copy and paste the certificate into the console/Telnet session, followed by a blank line to end the command.

History

This command was first available in the ExtremeXOS 11.2 and supported with the SSH module.

As of ExtremeXOS 21.1, the SSH XMOD is part of the base image and not available as a separate XMOD module.

Ability to configure CSR-signed certificates was added in ExtremeXOS 31.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ssl certificate privkeylen

```
configure ssl certificate privkeylen length country code organization  
                                org_name common-name name
```

Description

Creates a self-signed certificate and private key that can be saved in the EEPROM.

Syntax Description

<i>length</i>	Specifies the private key length in bytes. Valid values are between 2,048 and 4,096.
<i>code</i>	Specifies the country code in 2-character form.
<i>org_name</i>	Specifies the organization name. The organization name can be up to 64 characters long.
<i>name</i>	Specifies the common name. The common name can be up to 64 characters long.

Default

N/A.

Usage Guidelines

This command creates a self signed certificate and private key that can be saved in the EEPROM. The certificate generated is in the PEM format.

Any existing certificate and private key is overwritten.

The size of the certificate depends on the RSA key length (*privkeylen*) and the length of the other parameters (country, organization name, and so forth) supplied by the user. For an RSA key length of 4,096, the certificate length is approximately 2 Kb, and the private key length is approximately 3 Kb.

Example

The following example creates an SSL certificate in the USA for a website called bigcats:

```
configure ssl certificate privkeylen 2048 country US organization IEEE common-name bigcats
```

History

This command was first available in the ExtremeXOS 11.2 and supported with the SSH module.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ssl csr

```
configure ssl csr privkeylen length country code organization org_name  
common-name name
```

Description

Generates certificate signing request (CSR) and private key.

Syntax Description

ssl	Specifies SSL (Secure Sockets Layer).
csr	Specifies creating a CSR (certificate signing request).
privkeylen	Specifies setting the private key length.
<i>length</i>	Specifies the value for the private key length in bytes (2,048–4,096).
country	Specifies setting the country code.
<i>code</i>	Specifies the two-character value for the country code.
organization	Specifies setting the organization name.
<i>org_name</i>	Specifies the value for the organization name (maximum of 64 characters).
common-name	Specifies setting the common name.
<i>name</i>	Specifies setting the value for the common name (maximum of 60 characters).

Default

N/A.

Usage Guidelines



Note

There can only be one CSR per switch.

After entering values for the private key length, country code, organization, and common name, you are prompted to enter information for the Distinguished Name (DN): state, locality, organization unit, and email address.



Note

Due to changes in the Distinguished Name (DN), you are prompted to provide country, organization, and common name to ensure backward compatibility.

Example

The following example creates a CSR with a private key length of 2,048, country is USA, organization is "EXTR", and the common name is "test":

```
# configure ssl csr privkeylen 2048 country US organization EXTR common-name test
You are about to be asked to enter information that will be incorporated into your
certificate request.
What you are about to enter is what is called a Distinguished Name or a DN.
For some fields there will be a default value in [].
If you enter '.' the field will be left blank.
-----
State or Province Name (full name) []: North Carolina
Locality Name (eg, city) [Default City]: Raleigh
Organizational Unit Name (eg, section) []: RDU
Email Address []: jsmith@extremenetworks.com
```

```

.....+++
.....+++
CSR and Key Pair generated.
-----BEGIN CERTIFICATE REQUEST-----
MIIC3TCCAcUCAQIwGZcxZcZAJBgNVBAYTA1VTMQ0wCwYDVQQKDARFWFRSMREwDwYD
VQODDAhjc3JfdGVzdDEXMBUGA1UECAwOTm9ydGggQ2Fyb2xpbmExEDAOBgNVBACM
B1JhbGVpZ2gxDDAKBgNVBAsMA1JEVTEtMCsGCSqGSIb3DQEJARYebHBldHR5am9o
bkBleHRyZW1lbmV0d29ya3MuY29tMIIBIjANBgkqhkiG9w0BAQEFAAOCAQ8AMIIB
CgKCAQEAm43c60n1XXkk1MMvK+ovX8fAhWRu8j7TAKGrSEnqEhmS0BI05bjZLsj/
loulgsFXQA17W4010OMt5w9zcMCNmSf47PJwpQZpo4msAW8uSp7IMM9Ctv0a8oLr
kArzh3F+Gp0cAe7LycOthiXINKKwmzWpNwHmGbrwAhbd3grShurvUU7n0b+1Xcle
YH5J/HnGq+j6Lb+iNF2RbCactChF0aeT7DKXZaIt8s+p9ib3XQXUNvGoP+4M/Eoq
dHfOwpvBJeL3EyhjkEmz456nwdtsY8deNi/ssW+VJJWpGPONNLo+11wD7BksCPtJ
Pf20atDCFj6bFAo6N9gbdkh1dI3euwIDAQABoAAwDQYJKoZIhvcNAQENBQADggEB
AikoEBWhrPmL4tf0KSgKeadfODJ6Nipkcyof9YZ9AceJhtgMmBfMmfcUrE+3e28j
asXQpEc5hLkc8fyRMNjDHuuz2d6uWju+K/TqVNT094bvbvySFsdBKjLcOAdlRP0m
CIMCCiAiaFhtmlE5Sg6BoYctJ2jRNJ4UQOejeclcg80+qaXu6u7xAg5emGMtJizE
bvePhgSdhYTCFGngFrg3pZXHHTvRB7t54oYGG7yYdFb3jyW8CzckxnkiTV87fxHP
ojUeAwXet1AfI8cof1Dfmf6gKnBLMzrz5DMDmqdJgE2HgLLZCLv+JZbjbmowLrDL
DhG3F97QQkwROTpJfmrSsaU=
-----END CERTIFICATE REQUEST-----

Warning: SSL Certificate and Key will not match now.
Please load new CA signed certificate.
New Key will be usable after restart of thttpd process.
Storing the private key. This may take some time.
.Done

```

History

This command was first available in ExtremeXOS 22.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure ssl privkey pregenerated

```
configure ssl privkey pregenerated
```

Description

Obtains the pre-generated private key from the user.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command is also used when downloading or uploading the configuration. The private key is stored in the EEPROM, and the certificate is stored in the configuration file.

With this command, you copy and paste the private key into the command line followed by a blank line to end the command. The following security algorithms are supported:

- RSA for public key cryptography (generation of certificate and public-private key pair, certificate signing). RSA key size between 1024 and 4096 bits.
- Symmetric ciphers (for data encryption): RC4, DES, and 3DES.
- Message Authentication Code (MAC) algorithms: RSA Data Security, Inc. [MD5](#) Message-Digest Algorithm and SHA.

The certificate and private key file should be in PEM format and generated using RSA as the cryptography algorithm.

Example

The following command obtains the pre-generated private key from the user:

```
configure ssl privkey pregenerated
```

Next, you open the certificate and then copy and paste the certificate into the console/Telnet session, followed by **[Enter]** to end the command.

History

This command was first available in the ExtremeXOS 11.2 and supported with the SSH module.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure stack-ports debounce time

```
configure stack-ports {port-list} debounce time [default | time]
```

Description

Configures debounce time feature on stacking ports.

Syntax Description

<i>port-list</i>	Specifies one or more stacking ports.
default	Configure the default value.
milliseconds	Time in milliseconds. Range is 0 (no debouncing) to 5000.

Default

Default debounce time value is 0.

Usage Guidelines

Debounce timer can be configured to override the false link flaps i.e. link flaps that happens in a milliseconds interval.

Example

```
configure stack-ports 1:1 1:2 debounce time 150
```

History

This command was first available in ExtremeXOS 15.3.4.

Platform Availability

The command is available on all stackable switches.

configure stacking alternate-ip-address

```
configure stacking alternate-ip-address [ipaddress netmask | ipNetmask]  
gateway automatic configure stacking [node-address node-address  
| slot slot_number] alternate-ip-address [ipaddress netmask |  
ipNetmask] gateway
```

Description

Configures an alternate management IP address, subnetwork, and gateway.

Syntax Description

<i>node-address</i>	Specifies the MAC address of a node in the stack. To view the MAC addresses for all nodes in a stack, enter the show stacking command. A node address or slot number is required unless the automatic keyword is specified.
<i>slot_number</i>	Specifies the slot number of the target node. To view the slot numbers, enter the show stacking command.
<i>ipaddress netmask</i>	Specifies the unique address that exists on the Management VLAN subnet as configured on the initial master node together with the subnetwork mask specified for the Management subnetwork. Example: 66.77.88.1 255.255.255.0.

<i>ipNetmask</i>	Specifies the unique address that exists on the Management VLAN subnet as configured on the initial master node, followed by a slash (/) character, followed by a decimal number that represents the number of leading one bits in the subnetwork address. An example is 66.77.88.1/24.
<i>gateway</i>	The address of an IP router. A default route is set up to reach this gateway.

Default

No alternate IP address is configured.

Usage Guidelines

If a Management subnetwork is configured and the alternate IP subnetwork does not exactly match the configured Management subnetwork, the information configured by one of the commands specified above is not used. The previously configured alternate IP address is removed if it was installed and subsequently a Management subnetwork is configured that does not exactly match the alternate IP subnetwork. In either case, an error message is logged. The alternate IP address is used if there is no configured Management subnetwork.

To use the command with the node address, the node must be in the stack topology; and to use the command with the slot number, the node must be in the active topology. This form of the command operates only on one node at a time. There are no checks to verify that the address is the one configured in the management VLAN subnet.

The command that does not require a node address or slot number specifies the automatic keyword. Usage of this form of the command causes an alternate IP address to be assigned to every node in the stack topology. The first address is the address specified in the [*ipaddress netmask* | *ipNetmask*] parameter. The next address is the IP address plus one, and so on. Since there is a specified subnet mask, the address is checked to insure that the block of IP addresses fits within the specified subnet given the number of nodes in the stack topology. The range of addresses is tested to insure that each one is a valid IP unicast address. If the test fails, no node is configured and an error message is printed. Assignment is in the order in which nodes would currently appear in the [show stacking](#) display.

The configuration takes effect after the command is successfully executed.

The alternate IP address, subnetwork, and gateway are only used when the node is operating in stacking mode.

Example

To configure an alternate IP address for every node in the stack with a single command:

```
configure stacking alternate-ip-address 10.120.1.10/24 10.120.1.1 automatic
```

To configure an alternate IP address on a single node in the stack topology:

```
configure stacking node-address 00:04:96:26:6b:ed alternate-ip-address 10.120.1.1/24 10.120.1.1
```

You may configure an alternate IP address using a slot number for a node that is currently occupying the related slot:

```
configure stacking slot 4 alternate-ip-address 10.120.1.13/24 10.120.1.1
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available with all licenses and platforms that support the SummitStack feature. For information about which licenses and platforms support the SummitStack feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure stacking easy-setup

```
configure stacking easy-setup
```

Description

This command provides an easy way to initially configure the stacking parameters of all nodes in a new stack.

Syntax Description

This command does not have additional syntax.

Default

N/A.

Usage Guidelines

This command performs the following functions:

- Informs you of the stacking parameters that will be set.
- Informs you of the number of nodes that will be configured.
- Informs you whether minimal or no redundancy will be configured, and which slot will contain the master node.
- Informs you of the slot number that will be assigned to the node on which your management session is being run.
- If applicable, warns you that the current configuration file changes will be lost and you need to save the files.
- If the stack topology is a daisy chain, warns you that you should wire the stack as a ring before running this command.

- Requires you to confirm before the operation takes place. If you proceed, the command does the following:
 - Enables stacking on all nodes.
 - Configures the stacking MAC address using the factory address of the current node.
 - Configures a slot number for each node.
 - Configures redundancy to minimal in a ring topology or none in a daisy chain topology.
 - Configures the stacking protocol.
 - Reboots the stack topology.
- Selects the enhanced stacking protocol.

Stacking is enabled as if the `enable stacking {node-address node-address}` command was issued.

The stack mac-address is configured as if the `configure stacking mac-address` was issued on the current node.

Stack slot numbers are assigned as if the `configure stacking slot-number automatic` command was issued on the current node.

On a daisy chain topology, the master-capability is configured as if the `configure stacking redundancy none` command was issued. On a ring topology, the master-capability is configured as if the `configure stacking redundancy minimal` command was issued.

If you choose not to proceed with the setup, the following message is displayed:

Cancelled easy stack setup configuration.

Example

If you have an 8-node stack in a ring topology and have powered on all the nodes, the `show stacking` command shows the stack topology as a ring with all intended nodes present. If you have not changed any ExtremeXOS configuration, the command displays as follows:

```
* Switch.30 # configure stacking easy-setup
For every node in the 8-node stack, this command will:
- enable stacking
- configure a stack MAC address
- choose and configure a slot number (this node will be assigned to slot 1)
- configure redundancy to minimal (slot 1 will be the Master node)
Upon completion, the stack will automatically be rebooted into the new configuration.
Warning: If stacking is already configured, this command will alter that configuration.
Warning: There are unsaved configuration changes. You may wish to save them before
proceeding.
Do you wish to proceed? (y/N) y
Stacking configuration is complete. Rebooting...
```

If the 8-node stack topology is a daisy chain, and the user is logged into a node in the middle of the chain, the command output might appear as follows:

```
* Switch.30 # configure stacking easy-setup
For every node in the 8-node stack, this command will:
- enable stacking
- configure a stack MAC address
```

```
- choose and configure a slot number (this node will be assigned to slot 5)
- configure redundancy to none (slot 1 will be the master node)
Upon completion, the stack will automatically be rebooted into the new configuration.
Warning: If stacking is already configured, this command will alter that configuration.
Warning: This stack is a daisy chain. It is highly recommended that the stack
be connected as a ring before running this command.
Do you wish to proceed? (y/N) Yes
Stacking configuration is complete. Rebooting...
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available with all licenses and platforms that support the SummitStack feature. For information about which licenses and platforms support the SummitStack feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure stacking license-level

```
configure stacking {node-address node-address | slot slot-number}
license-level license_restriction
```

Description

Allows you to restrict the license level at which the node operates.

Syntax Description

<i>node-address</i>	Specifies the MAC address of a node in the stack. To view the MAC addresses for all nodes in a stack, enter the show stacking command.
<i>slot-number</i>	Specifies the slot number of the target node. To view the slot numbers, enter the show stacking command.
<i>license_restriction</i>	Specifies the restricted license level: <ul style="list-style-type: none"> For ExtremeSwitching 5420 and 5520 series switches, the choices are Base or Premier.

Default

No license level restriction is configured.

Usage Guidelines

This command causes a node to operate at a lower license level than the level that was purchased for the node.

Running this command does not change the installed license level. For example, if a stackable is configured with the Advanced Edge license and you configure a license level restriction of Edge, the unit is restricted to features available in the Edge license. However, you can remove the restriction and operate at the Advanced Edge level.

If the installed license level of the target node is lower than the level you are attempting to configure, the following message appears:

Warning: Switch will not operate at a license level beyond that which was purchased.

If the node-address or slot parameter is not specified, the command takes effect on every node in the stack topology.

This command takes effect after you restart the node. The following message appears after the command is executed:

This command will take effect at the next reboot of the specified node(s).

If you restart the node without configuring a license level restriction, the node operates at the purchased license level. To see the purchased license level of a node, run `show licenses` after logging in to the node.

The `show licenses` command displays the current license level in use as the Effective License Level:

```
Slot-2 Stack.1 # show licenses
Enabled License Level:
Advanced Edge
Enabled Feature Packs:
None
Effective License Level:
Edge
```

The `show stacking configuration` and `show stacking {node-address node-address | slotslot-number} detail` commands allow you to see the configured license level restriction and the restriction currently in use.

The Effective License Level appears only when stacking is enabled. The command is node-specific. The effective license level is the level at which the node is restricted to operate, and is not necessarily the level at which the entire stack is operating. This is because it is possible to have the restriction differ on each node, in which case one or more nodes may have failed because of the differing levels.

Example

To configure the stacking level Edge on all nodes in a stack:

```
# configure stacking license-level edge
```

To configure stacking level Edge for a node:

```
# configure stacking node-address 00:04:96:26:6b:ed license-level edge
```

To configure the stacking level Advanced Edge for an active node that currently occupies slot 4:

```
# configure stacking slot 4 license-level advanced-edge
```

History

This command was first available in ExtremeXOS 12.0.

The *license_restriction* variable was added, and the options **Edge**, **Advanced Edge**, and **Core** were removed in ExtremeXOS 31.1.

Platform Availability

This command is available with all licenses and platforms that support the SummitStack feature. For information about which licenses and platforms support the SummitStack feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure stacking mac-address

```
configure stacking {node-address node-address | slot slot-number} mac-  
address
```

Description

Selects a node in the stack whose factory assigned MAC address is to be used to form the stack MAC address.

The formed address is then configured on every node in the stack topology.

Syntax Description

<i>node-address</i>	Specifies the MAC address of a node in the stack. To view the MAC addresses for all nodes in a stack, enter the show stacking command.
<i>slot-number</i>	Specifies the slot number of an active node whose factory MAC address is to be used to form the stack MAC address. To view the slot numbers, enter the show stacking command.

Default

No stack MAC selection is configured.

Usage Guidelines

You must select a node whose factory assigned MAC address can be used to form a MAC address that represents the stack as a whole. The system forms the stack MAC address by setting the Universal / Local bit in the specified MAC address. This means that the stack MAC address is a locally administered address, and not the universal MAC address assigned to the selected node.

If you do not specify any node, the stack MAC address is formed from the factory assigned MAC address of the node from which you are running the command.

This command takes effect only after you restart the node. The following message appears after you run the command:

This command will take effect at the next reboot of the specified node(s).

If a stack node that has just joined the active topology detects that its stack MAC address is not configured or is different than the stack MAC address in use, it logs the following message at the Error log level:

```
The stack MAC address is not correctly configured on this node. The
stack can not operate properly in this condition. Please correct and
reboot.
```

If you have not configured (or inconsistently configured) the stack MAC address you might encounter difficulty in diagnosing the resulting problems. Whenever the master node (including itself) detects that one or more nodes in its active topology do not have the correct or any stack MAC address configured, it displays the following message to the console every five minutes until you configure a MAC address and restart the node(s):

```
The stack MAC address is either not configured or its configuration is
not consistent within the stack. The stack can not operate properly in
this condition. Please correct and reboot.
```

Example

To select the node to which you have logged in to supply the MAC address for stack MAC address formation:

```
configure stacking mac-address
```

To select a node other than the one to which you are logged in to supply the MAC address for stack MAC address formation:

```
configure stacking node-address 00:04:96:26:6b:ed mac-address
```

To select an active node to supply the MAC address for stack MAC address formation:

```
configure stacking slot 4 mac-address
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available with all licenses and platforms that support the SummitStack feature. For information about which licenses and platforms support the SummitStack feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure stacking master-capability

```
configure stacking [node-address node_address | slot slot_number]
master-capability [on | off]
```

Description

The command configures a node to be allowed to operate as either a backup or master, or prevents a node from operating as either.

The command controls the setting on the specified node only. To set the master capability for all nodes on a stack, you can use the command `configure stacking redundancy [none | minimal | maximal]`.

Syntax Description

<i>node_address</i>	Specifies the MAC address of a node in the stack. To view the MAC addresses for all nodes in a stack, enter the <code>show stacking</code> command.
<i>slot_number</i>	Specifies the slot number of the target active node. To view the slot numbers, enter the <code>show stacking</code> command.

Default

Master-capability is On.

Usage Guidelines

At least one node in the stack topology must be master-capable.

If you attempt to disable the master-capability of the only master capable node in a stack topology, the attempt is denied and following message appears:

```
Error: At least one node must have Master-capability configured "on".
```

This command is used to set up master-capability manually. It can also be used to adjust the result achieved when the `configure stacking redundancy [none | minimal | maximal]` command is used.

The setting takes effect the next time the node reboots. When this command is executed successfully, the following message appears:

```
This command will take effect at the next reboot of the specified
node(s) .
```

Example

To turn on the master capability for a node:

```
configure stacking node-address 00:04:96:26:6b:ed master-capability on
```


To turn on the master capability of an active node currently occupying slot 4:

```
configure stacking slot 4 master-capability on
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available with all licenses and platforms that support the SummitStack feature. For information about which licenses and platforms support the SummitStack feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure stacking node-address

```
configure stacking node-address node_address slot-number slot_number
```

Description

Configures a slot number on one or all nodes in the stack topology.

Syntax Description

<i>node_address</i>	Specifies the MAC address of a node in the stack. To view the MAC addresses for all nodes in a stack, enter the show stacking command.
<i>slot_number</i>	Specifies a number between 1 and 8 that is to be assigned as the slot number of the target node.

Default

The default slot-number for a node in stacking mode is 1.

Usage Guidelines

The configuration is stored on the affected node(s) immediately but does not take effect until the next reboot of the node(s). The configuration applies only when the node is running in stacking mode. To see the configured and active slot numbers of all nodes, use the [show stacking configuration](#) command.

If a node-address and a slot number are specified, then the node is configured with the specified slot number. There is no check for a duplicate slot number at this time; the number is simply assigned as requested.

To see the resulting slot number assignment, run the `show stacking configuration` command.



Note
Failure to configure a node does not prevent configuration of the slot numbers on the other nodes, and does not affect the slot number assigned to each node.

When this command is executed successfully, the following message appears:
`This command will take effect at the next reboot of the specified node(s).`

Example

To configure slot number 4 for the node with MAC address 00:04:96:26:6b:ed:

```
configure stacking node-address 00:04:96:26:6b:ed slot-number 4
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available with all licenses and platforms that support the SummitStack feature. For information about which licenses and platforms support the SummitStack feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure stacking priority

```
configure stacking {node-address node-address | slot slot_number}  
                  priority [node_pri | automatic]
```

Description

Configures a priority value to be used to influence master and backup election.

Syntax Description

<i>node-address</i>	Specifies the MAC address of a node in the stack. To view the MAC addresses for all nodes in a stack, enter the <code>show stacking</code> command.
<i>slot_number</i>	Specifies the slot number of the target node. To view the slot numbers, enter the <code>show stacking</code> command.
<i>node_pri</i>	Specifies the priority as a value between 1 and 100.

Default

Automatic priority.

Usage Guidelines

The node role election priority is a value that is internally calculated by ExtremeXOS for each node. This calculated value helps determine which nodes are elected as master and backup. For more information, see “Configuring the Master, Backup, and Standby Roles” in the [Switch Engine 32.3 User Guide](#).

This command allows you to configure a priority value that affects the outcome of this calculation. You can configure the priority on any node in a stack topology. You can specify an integer node-pri value between 1 and 100. The larger the value, the greater the node role election priority.

If no node address or slot is specified, the command takes effect on all nodes at the next node role election cycle. Priority configuration has no operational effect on switches that are not in stacking mode.

If configured on every node, automatic priority commands ExtremeXOS to determine the node role election priority of each active node. Currently, the automatic priority algorithm chooses the master-capable node with the lowest slot number as master and the node with the second lowest slot number as backup. Extreme networks may alter this behavior in later releases.

If you have configured a node with automatic priority and if you have configured another node to use a node-pri value, the node with automatic priority uses zero as the node-priority value during the node role election.

Example

To allow ExtremeXOS to determine node role election priority:

```
configure stacking priority automatic
```

To configure the node priority for the stackable in slot 4:

```
configure stacking slot 4 priority 50
```

To configure the automatic priority algorithm for the stackable with node address 00:04:96:26:6b:ed:

```
configure stacking node-address 00:04:96:26:6b:ed priority automatic
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available with all licenses and platforms that support the SummitStack feature. For information about which licenses and platforms support the SummitStack feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure stacking redundancy

```
configure stacking redundancy [none | minimal | maximal]
```

Description

This command sets a master-capability value for every node in the stack topology.

Syntax Description

none	Only one node has master-capability turned on and all other nodes have master-capability turned off.
minimal	Two nodes have master-capability turned on and all other nodes have master-capability turned off.
maximal	All nodes have master-capability turned on.

Default

Default value in an unconfigured stack is maximal.

Usage Guidelines

If there are more than eight nodes in the stack topology, the following message appears and the command is not executed:

```
ERROR: This command can only be used when the stack has eight nodes or less.
```

Since only eight nodes can be operational in an active topology at a time, you must disconnect the remaining nodes before configuring master-capability with this command.

If you are using the none or minimal redundancy configuration:

- The configured values of slot-number and priority decide the nodes on which the master-capability should be turned on.
- If the priority values are configured on the nodes, the highest priority node(s) are chosen.
- If the priority values of all nodes are set to automatic or to the same priority value, the node(s) with the lowest slot number(s) are chosen. Extreme Networks may change automatic priority behavior in a future release.

If there is a slot number tie or if the slot numbers were never configured, the following message appears and the command is not executed:

```
ERROR: Unique slot numbers must be configured before using this command.
```

The setting takes effect at the next restart of the node. The following message appears after the command is successfully executed:

```
This command will take effect at the next reboot of the specified node(s).
```

Redundancy configuration has no operational effect on a node that is not in stacking mode.

Example

To turn on master-capability on all nodes:

```
configure stacking redundancy maximal
```

To turn on master-capability on only one node:

```
configure stacking redundancy none
```

To turn on master-capability on two nodes:

```
configure stacking redundancy minimal
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available with all licenses and platforms that support the SummitStack feature. For information about which licenses and platforms support the SummitStack feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure stacking slot-number automatic

```
configure stacking slot-number automatic
```

Description

Configures a slot number on all nodes in the stack topology, selecting the number automatically.

Syntax Description

automatic	Configures slot numbers on every node in the stack, selecting the number automatically. If there are more than eight nodes in the stack topology, the assignment is only performed on the first eight nodes. Automatic slot number assignment causes assignment of slot numbers starting from 1 and increasing up to 8. The nodes in the stack topology are assigned the numbers in the order in which they would appear currently in the show stacking command output. In a ring, slot number 1 is assigned to the current node, slot number 2 is assigned to the node connected to the current node's stack port 2, and so forth. In a daisy chain, slot 1 is assigned to the node at the end of the chain that begins with the node connected to the current node's stack port 1.
------------------	---

Default

The default slot-number for a node in stacking mode is 1.

Usage Guidelines

The configuration is stored on the affected node(s) immediately but does not take effect until the next reboot of the node(s). The configuration applies only when the node is running in stacking mode. To see the configured and active slot numbers of all nodes, use the `show stacking configuration` command.

To see the resulting slot number assignment, run the `show stacking configuration` command.



Note

Failure to configure a node does not prevent configuration of the slot numbers on the other nodes, and does not affect the slot number assigned to each node.

If you enter the command with the automatic option, the following confirmation message appears:
Reassignment of slot numbers may make the stack incompatible with the current configuration file. Do you wish to continue? (y/n)

When this command is executed successfully, the following message appears:

This command will take effect at the next reboot of the specified node(s).

Example

To configure all slot-numbers for a stack:

```
configure stacking slot-number automatic
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available with all licenses and platforms that support the SummitStack feature. For information about which licenses and platforms support the SummitStack feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure stacking-support auto-discovery

```
configure stacking-support auto-discovery [disable | enable]
```

Description

Enables or disables stacking auto-discovery.

Syntax Description

stacking-support	Configures stacking support.
auto-discovery	Configures auto-discovery for stacking.
disable	Disables stacking auto-discovery.
enable	Enables stacking auto-discovery. (Default)

Default

Stacking auto-discovery is enabled by default.

Usage Guidelines

For ExtremeSwitching 5420 and 5520 series switches, stacking auto-discovery allows the switch to detect the type of cable inserted into the stack ports and automatically update the stack port speed while the switch is booting up.

To view stacking auto-discovery status, use the command `show stacking-support`.

Example

The following example disables stacking auto-discovery:

```
# configure stacking-support auto-discovery disable
```

History

This command was first available in ExtremeXOS 30.3.

Platform Availability

This command is available on ExtremeSwitching 5420 and 5520 series switches.

configure stacking-support stack-ports

```
configure stacking-support stack-port [stack-ports | all] selection  
    [native {V40 | V80 | V160 | V200 | V320 | V400 {alternative-  
        configuration | help}} | alternate]
```

Description

Selects the switch ports and speed for stack communications.

Syntax Description

<i>stack-ports</i>	Specifies the stacking port range to be configured. Valid stacking port entries are 1, 2, 1-2, and all.
native	Selects the specified stacking port, which is the native, dedicated port that only supports stacking.
v40	Specifies that the native stacking port can operate at 10 Gbps.
v80	Specifies that the native stacking port can operate at 20 Gbps.
v160	Specifies that the native stacking ports operate at 40 Gbps.
v200	Specifies that the native stacking ports operate at 50 Gbps
v320	Specifies that the native stacking ports operate at 80 Gbps.
v400	Specifies that the native stacking ports on the option card operate at 100 Gbps (Not available on Universal platforms).
alternative-configuration	Selects the V400 alternate configuration stacking mode for (ExtremeSwitchin 5720 series switches).
help	Provides more details regarding the alternate configuration stacking mode
alternate	Selects the alternate (Ethernet) stacking port associated with the specified stacking port. The alternate port numbers are listed in the following table.

Default

Switches with native stack ports default to "Native".

Usage Guidelines

The configuration entered with this command applies to only the local node and does not become active until after the following events:

- The stacking-support option is enabled (if applicable).
- The switch restarts.

Each speed configuration requires a specific cabling configuration. For more information, see *Hardware Installation Guide* for your switch model.

"V400" is the default mode that sets the stack ports to 106G. "V400 alternative-configuration" is required when using specific fiber cables. This mode sets the stack ports to 100G, enables pre-emphasis, and FEC (clause_91). Cables requiring alternative-configuration include:

- QSFP28 SR4
- QSFP28 LR4
- QSFP28 CWDM4
- QSFP28 PSM4
- QxQ AOC cable - 5m
- QxQ AOC cable - 7m
- QxQ AOC cable - 10m
- QxQ AOC cable - 20m

For a complete list of supported cables, see [Extreme Optics Compatibility](#).

The stacking-support option configures the switch to use stacking protocols. This option is automatically enabled on most platforms, but some platforms require you to manually enable the stacking-support option. The following table lists the ExtremeSwitching series switches and option card configurations that support Stacking Port Selection Control, and it lists which platforms require manual Stacking-Support Option Control.

Table 18: ExtremeSwitching Series Switch Support for Alternate Stack Ports

Switch Model Number	Switch Option Card	Alternate Port for Stack Port	Alternate Port for Stack Port	Stacking-Support Option Control	Stacking Port Selection Control
5520-24T 5520-24W 5520-24X	5520-VIM-4X	35	36	Yes	Yes
5520-48T 5520-48W 5520-12MW-36W 5520-48SE	5520-VIM-4X	59	60	Yes	Yes

When the alternate stack port is selected for a native stack port and the switch is restarted, the native stack port remains visible in the CLI and can be configured. However, any configuration applied to the replaced stack port is ignored and does not affect switch operation.

An alternate stack port runs the stacking protocol and cannot operate on a link connected to a data port that is not configured as a stack port. Both ends of a stack link must be configured to use the stacking protocol. The stacking link must be directly connected to two the alternate stacking ports of two stacking switches. The direct connection is necessary because stacking protocols cannot pass through an intermediate switch.

After a data port is reconfigured as an alternate stack port, all data port configuration commands still work, but they do not change the operation of the alternate stack port. The LEDs on an Ethernet port used as an alternate stacking port operate according to the behavior of the Ethernet port. The LEDs on the related (disabled) native stacking port remain dark.



Note

Commands that contain the **stacking-support** keyword operate only on the local switch; they do not apply to all switches in the stack. If an active stack topology has been formed, you can telnet to a slot elsewhere in the stack, log on to that switch, and use commands with the stacking-support keyword on that switch.

Example

The following command configures the switch to use the alternate stack port for Stack Port 1 after the next switch restart:

```
configure stacking-support stack-ports 1 selection alternate
```

The following command configures the switch to use both native stacking ports after the next switch restart:

```
configure stacking-support stack-ports all selection native
```

The following command configures stack ports 1 and 2 to operate as four 40 Gbps ports:

```
configure stacking-support stack-ports all selection native V160
```

History

This command was first available in ExtremeXOS 12.5.

The V160 keyword was added in ExtremeXOS 12.6.

The V320 keyword was added in ExtremeXOS 15.1 Revision 2.

The V400 keyword was added in ExtremeXOS 22.2.

The **alternative-configuration** and **help** keywords were added in ExtremeXOS 32.3.

The V40 keyword was added in ExtremeXOS 31.3

The V80 keyword was added in ExtremeXOS 31.4.

The V200 keyword was added in ExtremeXOS 31.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure stpd add vlan

```
configure stpd stpd_name add [ {vlan} vlan_name | vlan vlan_list] ports
    [all | port_list] {[dot1d | emistp | pvst-plus]}
```

Description

Adds all ports or a list of ports within a VLAN to a specified STPD.

Syntax Description

<i>stpd_name</i>	Specifies an STPD name on the switch.
<i>vlan_name</i>	Specifies a VLAN name.
<i>vlan_list</i>	Specifies a VLAN list of IDs.
all	Specifies all of the ports in the VLAN to be included in the STPD.
<i>port_list</i>	Specifies the port or ports to be included in the STPD.
dot1d	Specifies the <u>STP</u> encapsulation mode of operation to be 802.1D.

emistp	Specifies the STP encapsulation mode of operation to be EMISTP.
pvst-plus	Specifies the STP encapsulation mode of operation to be PVST+.

Default

Default port mode for default STPD (s0) and user-created STPDs is dot1d.

Usage Guidelines

To create an STP domain, use the `create stpd` command. To create a VLAN, use the `create vlan` command.

In an EMISTP or PVST+ environment, this command adds a list of ports within a VLAN to a specified STPD provided the carrier VLAN already exists on the same set of ports. You can also specify the encapsulation mode for those ports.

In an *MSTP* environment, you do not need a carrier VLAN. A CIST controls the connectivity of interconnecting MSTP regions and sends BPDUs across the regions to communicate region status. You must use the dot1d encapsulation mode in an MSTP environment.

You cannot configure STP on the following ports:

- Mirroring target ports.
- Software-controlled redundant ports.

If you see an error similar to the following:

```
Error: Cannot add VLAN default port 3:5 to STP domain
```

You might be attempting to add:

- A carrier VLAN port to a different STP domain than the carrier VLAN belongs.
- A VLAN/port for which the carrier VLAN does not yet belong.



Note

This restriction is enforced only in an active STP domain and when you enable STP to make sure you have a legal STP configuration.

Care must be taken to ensure that ports in overlapping domains do not interfere with the orderly working of each domain's protocol.

By default, when the switch boots for the first time, it automatically creates a VLAN named default with a tag value of 1 and STPD s0. The switch associates VLAN default to STPD s0. All ports that belong to this VLAN and STPD are in 802.1D encapsulation mode with autobind enabled. If you disable autobind on the VLAN default, that configuration is saved across a reboot.

Naming Conventions

If your STPD has the same name as another component, for example a VLAN, we recommend that you specify the identifying keyword as well as the name. If your STPD has a name unique only to that STPD, the keywords **stpd** and **vlan** are optional.

STP Encapsulations Modes

You can specify the following STP encapsulation modes:

- dot1d—This mode is reserved for backward compatibility with previous STP versions. BPDUs are sent untagged in 802.1D mode. Because of this, any given physical interface can have only one STPD running in 802.1D mode.

This encapsulation mode supports the following STPD modes of operation: 802.1D, 802.1w, and MSTP.

- emistp—This mode sends BPDUs with an 802.1Q tag having an STPD ID in the VLAN ID field.

This encapsulation mode supports the following STPD modes of operation: 802.1D and 802.1w.

- 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.

This encapsulation mode supports the following STPD modes of operation: 802.1D and 802.1w.

These encapsulation 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.

MSTP STPDs use 802.1D BPDU encapsulation mode by default. To ensure correct operation of your MSTP STPDs, do not configure EMISTP or PVST+ encapsulation mode for MSTP STPDs.

STPD Identifier

An StpdID is used to identify each STP domain. You assign the StpdID when configuring the domain. An STPD ID must be identical to the VLAN ID of the carrier VLAN in that STPD and that VLAN cannot belong to another STPD.

MSTP uses two different methods to identify the STPDs that are part of the MSTP network. An instance ID of 0 identifies the Common and Internal Spanning Tree (CIST). The switch assigns this ID automatically when you configure the CIST STPD. A multiple spanning tree instance identifier identifies each STP domain that is part of an MSTP region. You assign the MSTI (Multiple Spanning Tree Instances) ID when configuring the STPD that participates in the MSTP region. In an MSTP region, MSTI IDs only have local significance. You can reuse MSTI IDs across MSTP regions.

Automatically Inheriting Ports--MSTP Only

In an MSTP environment, whether you manually or automatically bind a port to an MSTI in an MSTP region, the switch automatically binds that port to the CIST. The CIST handles BPDU processing for itself and all of the MSTIs; therefore, the CIST must inherit ports from the MSTIs in order to transmit and receive BPDUs.

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, and includes all the ports of the VLAN in STPD1:

```
configure stpd stpd1 add vlan marketing ports all
```

History

This command was first available in ExtremeXOS 10.1.

The `vlan_list` variable was added in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure stpd backup-root

```
configure stpd stpd_name backup-root [on | off]
```

Description

Enables and disables the backup root feature.

Syntax Description

<i>stpd_name</i>	Specifies an <i>STPD</i> name on the switch.
on	Enables backup root.
off	Disable backup root.

Default

By default, the backup root feature is disabled.

Usage Guidelines

The backup root feature is used to get faster convergence when the root bridge connectivity is lost.

Backup root feature enabled bridge port should be connected to Root with point to point link. When backup root bridge loses contact with the root bridge, the backup root bridge automatically lowers its bridge priority below the priority of the lost root. This causes the backup root bridge to become the new root. If a reboot occurs, the new root will have its priority restored to the original configured value.

If the priority of the root bridge is zero and the backup root loses connectivity to the root bridge, automatic assignment of the priority value for the backup root will be the initial configured value.

This feature is activated only when connectivity with the root bridge is lost. Raising the priority on the root does not cause the backup root feature to be activated.

We recommend the following when configuring the backup root feature:

- Enable the backup root feature on both the root and backup root.
- Configure all bridges except the root and backup root with the maximum bridge priority value (61440 with 802.1t).
- Configure the root and backup root to have the next lowest priority (57344 with 802.1t)
- To help prevent the backup root feature activating due to a simple link failure rather than a bridge failure, establish multiple links between the root and backup root.
- Deploy this feature carefully as it may result in suboptimal traffic forwarding paths.

Example

The following example enables the backup root feature on the *STP* domain r1:

```
configure stpd r1 backup-root on
```

History

This command was first available in ExtremeXOS 15.7.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure stpd bpdud-forwarding

```
configure stpd bpdud-forwarding [on | off]
```

Description

This command specifies whether to forward or drop BPDUs when *STP* is disabled.

Syntax Description

bpdud-forwarding	Specifies forwarding or discarding spanning tree BPDUs when STP is disabled.
on	Forward STP BPDUs when spanning tree is disabled (default).
off	Drop STP BPDUs when spanning tree is disabled.

Default

The default is on.

Usage Guidelines

STP must be disabled globally to disable BPDU forwarding; otherwise, an error message appears:
Error: All Spanning Tree Domains must be disabled globally before configuring stpd bpdud-forwarding off.

When the BPDU forwarding is off and you try to configure the filter method using the `configure stpd filter-method` [**system-wide** | **port-based**] command, the following error message appears:

Error: Spanning Tree Forwarding must be enabled globally before configuring filter-method.

Example

The following example disables BPDU forwarding when STP is disabled:

```
configure stpd bpdu-forwarding off
```

History

This command was first available in ExtremeXOS 22.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure stpd default-encapsulation

```
configure stpd stpd_name default-encapsulation [dot1d | emistp | pvst-plus]
```

Description

Configures the default encapsulation mode for all ports added to the specified *STPD*.

Syntax Description

<i>stpd_name</i>	Specifies an STPD name on the switch.
dot1d	Specifies the <i>STP</i> encapsulation mode of operation to be 802.1d.
emistp	Specifies the STP encapsulation mode of operation to be EMISTP.
pvst-plus	Specifies the STP encapsulation mode of operation to be PVST+.

Default

Ports in the default STPD (s0) are dot1d mode.

Ports in user-created STPDs are in dot1d mode.

Usage Guidelines

Care must be taken to ensure that ports in overlapping domains do not interfere with the orderly working of each domain's protocol.

By default, when the switch boots for the first time, it automatically creates a VLAN named default with a tag value of 1 and STPD s0. The switch associates VLAN default to STPD s0. All ports that belong to this VLAN and STPD are in 802.1d encapsulation mode with autobind enabled. If you disable autobind on the VLAN default, that configuration is saved across a reboot.

MSTP STPDs use 802.1D BPDU encapsulation mode by default. To ensure correct operation of your MSTP STPDs, do not configure EMISTP or PVST+ encapsulation mode for MSTP STPDs.

Naming Conventions

If your STPD has the same name as another component, for example a VLAN, we recommend that you specify the identifying keyword as well as the name. If your STPD has a name unique only to that STPD, the keyword stpd is optional. For name creation guidelines and a list of reserved names, see *Object Name* in the [Switch Engine 32.3 User Guide](#).

STP Encapsulation Modes

You can specify the following STP encapsulation modes:

- dot1d—This mode is reserved for backward compatibility with previous STP versions. BPDUs are sent untagged in 802.1D mode. Because of this, any given physical interface can have only one STPD running in 802.1D mode.

This encapsulation mode supports the following STPD modes of operation: 802.1D, 802.1w, and MSTP.

- emistp—This mode sends BPDUs with an 802.1Q tag having an STPD ID in the VLAN ID field.

This encapsulation mode supports the following STPD modes of operation: 802.1D and 802.1w.

- 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.

This encapsulation mode supports the following STPD modes of operation: 802.1D and 802.1w.



Note

These encapsulation 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.

STPD Identifier

An StpdID is used to identify each STP domain. You assign the StpdID when configuring the domain. An STPD ID must be identical to the VLAN ID of the carrier VLAN in that STP domain, and that VLAN cannot belong to another STPD.

MSTP uses two different methods to identify the STPDs that are part of the MSTP network. An instance ID of 0 identifies the Common and Internal Spanning Tree (CIST). The switch assigns this ID automatically when you configure the CIST STPD. A multiple spanning tree instance identifier identifies each STP domain that is part of an MSTP region. You assign the MSTI ID when configuring the STPD that

participates in the MSTP region. In an MSTP region, MSTI IDs only have local significance. You can reuse MSTI IDs across MSTP regions.

Example

The following example specifies that all ports subsequently added to the STPD STPD1 be in PVST+ encapsulation mode unless otherwise specified or manually changed:

```
configure stpd stpd1 default-encapsulation pvst-plus
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure stpd delete vlan

```
configure stpd stpd_name delete [ {vlan} vlan_name | vlan
vlan_list]ports [all | port_list]
```

Description

Deletes one or more ports in the specified *VLAN* from an *STPD*.

Syntax Description

<i>stpd_name</i>	Specifies an STPD name on the switch.
<i>vlan_name</i>	Specifies a VLAN name.
<i>vlan_list</i>	Specifies a VLAN list of IDs.
all	Specifies that all of the ports in the VLAN are to be removed from the STPD.
<i>port_list</i>	Specifies the port or ports to be removed from the STPD.

Default

N/A.

Usage Guidelines

If your STPD has the same name as another component, for example a VLAN, we recommend that you specify the identifying keyword as well as the name. If your STPD has a name unique only to that STPD, the keywords stpd and vlan are optional.

In EMISTP and PVST+ environments, if the specified VLAN is the carrier VLAN, all protected VLANs on the same set of ports are also removed from the STPD.

You also use this command to remove autobind ports from a VLAN. ExtremeXOS records the deleted ports so that the ports are not automatically added to the STPD after a system restart.

When a port is deleted on the *MSTI*, it is automatically deleted on the CIST as well.

Example

The following example removes all ports of a VLAN named Marketing from the STPD STPD1:

```
configure stpd stpd1 delete vlan marketing ports all
```

History

This command was first available in ExtremeXOS 10.1.

The *vlan_list* variable was added in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure stpd description

```
configure {stpd} stpd_name description [stpd-description | none]
```

Description

Adds or overwrites the *STP* domain description field.

Syntax Description

<i>stpd_name</i>	Specifies an <i>STPD</i> name on the switch.
<i>stpd-description</i>	Specifies an STPD description.
none	Clears the STPD string.

Default

The STP domain description string is empty.

Usage Guidelines

Use this command to add or overwrite the STP domain description field.

The maximum STP domain description length is 180 characters.

The stpd-description must be in quotes if the string contains any spaces.

To display the description, use the `show stpd stpd_name` command. When no STP domain description is configured, Description is not displayed in the output.

To clear the STP domain description string, either specify the keyword `none` in this command or use the `unconfigure stpd {stpd_name}` command.

Example

The following command adds the description “this is s0 domain” to the STPD named s0:

```
configure stpd s0 description "this is s0 domain"
```

History

This command was first available in ExtremeXOS 12.4.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure stpd filter-method

```
configure stpd filter-method [system-wide | port-based]
```

Description

Configures Spanning Tree BPDU hardware filters.

Syntax Description

system-wide	Installs system-wide hardware filters for Spanning Tree.
port-based	Installs per-port hardware filters for Spanning Tree.

Default

By default, system-wide hardware filters are installed.

Usage Guidelines

You must disable Spanning Tree before changing the filter method. Use the `disable stpd` command to disable Spanning Tree.

Example

The following example sets the filter method for Spanning Tree as system-wide.

```
configure stpd filter-method system-wide
```

History

This command was first available in ExtremeXOS 15.7.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure stpd flush-method

```
configure stpd flush-method [vlan-and-port | port-only]
```

Description

Configures the method used by *STP* to flush the *FDB* during a topology change.

Syntax Description

vlan-and-port	Specifies a <i>VLAN</i> and port combination flush method.
port-only	Specifies a port flush method.

Default

The default flush method is vlan-and-port.

Usage Guidelines

For scaled up configurations where there are more than 1000 VLANs and more than 70 ports participating in STP, the number of messages exchanged between STP/FDB/HAL modules can consume a lot of system memory during an STP topology change using the default configuration for flush method. In such situations, setting the flush method to “port-only” can help reduce the system memory consumption.

Example

The following command sets the flush method to port-only:

```
configure stpd flush-method port-only
```

History

This command was available in ExtremeXOS 12.4.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure stpd forwarddelay

```
configure stpd stpd_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

<i>stpd_name</i>	Specifies an STPD name on the switch.
<i>seconds</i>	Specifies the forward delay time in seconds. The default is 15 seconds, and the range is 4 to 30 seconds.

Default

The default forward delay time is 15 seconds.

Usage Guidelines

If your STPD has the same name as another component, for example a *VLAN*, we recommend that you specify the identifying keyword as well as the name. If your STPD has a name unique only to that STPD, the keyword `stpd` is optional.

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:

```
configure stpd stpd1 forwarddelay 20
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure stpd hellotime

```
configure stpd stpd_name hellotime seconds
```

Description

Specifies the time delay (in seconds) between the transmission of BPDUs from this *STPD* when it is the root bridge.

Syntax Description

<i>stpd_name</i>	Specifies an STPD name on the switch.
<i>seconds</i>	Specifies the hello time in seconds. The default is 2 seconds, and the range is 1 to 10 seconds.

Default

The default hello time is 2 seconds.

Usage Guidelines

If your STPD has the same name as another component, for example a *VLAN*, we recommend that you specify the identifying keyword as well as the name. If your STPD has a name unique only to that STPD, the keyword stpd is optional.

In an *MSTP* environment, configure the hello timer only on the CIST, not on the MSTIs.

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:

```
configure stpd stpd1 hellotime 10
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure stpd loop-protect event-threshold

```
configure stpd stpd_name loop-protect event-threshold [threshold | none]
```

Description

Configures the loop protect event threshold.

Syntax Description

<i>stpd_name</i>	Specifies an <i>STPD</i> name on the switch.
<i>threshold</i>	Sets the number of loop protect events that must be received before disabling the port. The valid range is 1–10.
none	Disables the loop protect threshold. The port will not remain enabled even if loop protect events are received.

Default

By default, the loop protect threshold is enabled and set to three loop protect events.

Usage Guidelines

If the loop protect event threshold disables a port, you must enable the port manually.

Example

The following example configures the loop protect event threshold to five events.

```
configure stpd r1 loop-protect event-threshold 5
```

History

This command was first available in ExtremeXOS 15.7.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure stpd loop-protect event-window

```
configure stpd stpd_name loop-protect event-window interval
```

Description

Configures the interval for which loop protect events are counted by the loop protect event threshold.

Syntax Description

<i>stpd_name</i>	Specifies an <i>STPD</i> name on the switch.
<i>interval</i>	The length of the interval, in seconds, over which the loop protect event threshold is defined. The valid range is 0–255 seconds.

Default

By default the interval is set to 180 seconds.

Usage Guidelines

None.

Example

The following example sets the loop protect event window to 120 seconds for *STP* domain r1.

```
configure stpd r1 loop-protect event-window 120
```

History

This command was first available in ExtremeXOS 15.7.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure stpd maxage

```
configure stpd stpd_name maxage seconds
```

Description

Specifies the maximum age of a BPDU in the specified *STPD*.

Syntax Description

<i>stpd_name</i>	Specifies an STPD name on the switch.
<i>seconds</i>	Specifies the maxage time in seconds. The default is 20 seconds, and the range is 6 to 40 seconds.

Default

The default maximum age of a BPDU is 20 seconds.

Usage Guidelines

If your STPD has the same name as another component, for example a VLAN, we recommend that you specify the identifying keyword as well as the name. If your STPD has a name unique only to that STPD, the keyword `stpd` is optional.

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.

In an MSTP environment, configure the maximum age of a BPDU only on the CIST, not on the MSTIs.

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:

```
configure stpd stpd1 maxage 30
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure stpd max-hop-count

```
configure stpd stpd_name max-hop-count hopcount
```

Description

Specifies the maximum hop count of a BPDU until the BPDU is discarded in the specified MSTP STP domain.

Syntax Description

<i>stpd_name</i>	Specifies an <u>STPD</u> name on the switch.
<i>hopcount</i>	Specifies the number of hops required to age out information and notify changes in the topology. The default is 20 hops, and the range is 6 to 40 hops.

Default

The default hop count of a BPDU is 20 hops.

Usage Guidelines

This command is applicable only in an MSTP environment.

If your STPD has the same name as another component, for example a VLAN, Extreme Networks recommends that you specify the identifying keyword as well as the name. If your STPD has a name unique only to that STPD, the keyword `stpd` is optional.

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 `hopcount` parameter is 6 through 40 hops.

In an MSTP environment, the hop count has the same purpose as the maxage timer for 802.1D and 802.1w environments.

The main responsibility of the CIST is to exchange or propagate BPDUs across regions. The switch assigns the CIST an instance ID of 0, which allows the CIST to send BPDUs for itself in addition to all of the MSTIs within an MSTP region. Inside a region, the BPDUs contain CIST records and piggybacked M-records. The CIST records contain information about the CIST, and the M-records contain information about the MSTIs. Boundary ports only exchange CIST record BPDUs.

On boundary ports, only CIST record BPDUs are exchanged. In addition, if the other end is an 802.1D or 802.1w bridge, the maxage timer is used for interoperability between the protocols.

Example

The following command sets the hop of the MSTP STPD, STPD2, to 30 hops:

```
configure stpd stpd2 max-hop-count 30
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure stpd mode

```
configure stpd stpd_name mode [dot1d | dot1w | mstp [cist | msti  
instance]]
```

Description

Configures the operational mode for the specified STP domain.

Syntax Description

<i>stp_name</i>	Specifies an <i>STPD</i> name on the switch.
dot1d	Specifies the STPD mode of operation to be 802.1D.
dot1w	Specifies the STPD mode of operation to be 802.1w, and rapid configuration is enabled.
mstp	Specifies the STPD mode of operation to be 802.1s, and rapid configuration is enabled.
cist	Configures the specified STPD as the common instance spanning tree for the <i>MSTP</i> region.
msti	Configures the specified STPD as a multiple spanning tree instance for the MSTP region.
<i>instance</i>	Specifies the Id of the multiple spanning tree instance. The range is 1 to 4,094.

Default

The STPD s0 by default operates in MSTP CIST mode.

User-created STPDs operate by default in dot1d mode.

Usage Guidelines

If your STPD has the same name as another component, for example a *VLAN*, we recommend that you specify the identifying keyword as well as the name. If your STPD has a name unique only to that STPD, the keyword stpd is optional.

If you configure the STP domain in 802.1D mode, the rapid reconfiguration mechanism is disabled.

If you configure the STP domain in 802.1w mode, the rapid reconfiguration mechanism is enabled. You enable or disable RSTP on a per STPD basis only. You do not enable RSTP on a per port basis.

If you configure the STP domain in MSTP mode, the rapid reconfiguration mechanism is enabled. You enable or disable MSTP on a per STPD basis only. You do not enable MSTP on a per port basis. MSTP STPDs use 802.1D BPDU encapsulation mode by default. To ensure correct operation of your MSTP STPDs, do not configure EMISTP or PVST+ encapsulation mode for MSTP STPDs.

You must first configure a Common and Internal Spanning Tree (CIST) before configuring any multiple spanning tree instances (MSTIs) in the region. You cannot delete or disable a CIST if any of the MSTIs are active in the system.

STP operational mode can be changed while VLANs are associated with an STP domain. In MSTP mode, mode change is allowed only for CIST domains.

Example

The following command configures STPD s1 to enable the rapid reconfiguration mechanism and operate in 802.1w mode:

```
configure stpd s1 mode dot1w
```

The following command configures STPD s2 to operate as an *MSTI* in an MSTP domain:

```
configure stpd s2 mode mstp msti 3
```

History

This command was first available in ExtremeXOS 10.1.

The mstp parameter was added in ExtremeXOS 11.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure stpd multicast send-query

```
configure stpd multicast send-query [on | off]
```

Description

Configures suppressing *IGMP*- and MLD-triggered queries when *STP* topology changes are received.

Syntax Description

multicast	Specifies multicast options.
send-query	For VLANs associated with STPD, when topology changes occur, send or suppress IGMP or MLD queries.
on	Send IGMP or MLD queries (default).
off	Do not send IGMP or MLD queries.

Default

Sending IGMP or MLD queries is on.

Usage Guidelines

Whenever STP topology changes are received on a port, the switch sends triggered queries that mark the peer port as a router port and floods all multicast packets towards this port. This can cause unnecessary bandwidth usage. This command allows you to allow or suppress this forwarding.

Example

The following example turns off IGMP and MLD queries:

```
# configure stpd multicast send-query off
```

History

This command was first available in ExtremeXOS 21.1.5-Patch1-2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure stpd ports active-role disable

```
configure stpd stpd_name ports active-role disable port
```

Description

Allows a port to be selected as an alternate or backup port.

Syntax Description

<i>stpd_name</i>	Specifies an <i>STPD</i> name on the switch.
<i>port</i>	Specifies a port.

Default

The default is disabled.

Usage Guidelines

Use this command to revert to the default that allows a specified port to be elected to any *STP* port role.

Example

The following command disables an active role on STDP s1, port 6:3:

```
configure stpd s1 ports active-role disable 6:3
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure stpd ports active-role enable

```
configure stpd stpd_name ports active-role enable port
```

Description

Prevents a port from becoming an alternate or backup port.

Syntax Description

<i>stpd_name</i>	Specifies an <i>STPD</i> name on the switch.
<i>port</i>	Specifies a port.

Default

The default is disabled.

Usage Guidelines

Use this command to keep a port in an active role. It prevents a specified port from being elected to an alternate or backup role which puts the port in a blocking state.

The following describes the port role and state when RSTP stabilizes.

STP Port Role	Port State
Alternate (inactive)	Blocking
Backup (inactive)	Blocking
Root (active)	Forwarding
Designated (active)	Forwarding

This feature can be enabled on only one *STP* port in the STP domain.

The restricted port role cannot be combined with this feature.

An active port role (root or designated) cannot be enabled with an edge port.

To disable this command, use the `configure stpd ports active-role disable` command.

To view the status of the active role, use the `show stpd ports` command.

Example

The following command enables an active role on STDP s1, port 6:3:

```
configure stpd s1 ports active-role enable 6:3
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure stpd ports auto-edge

```
configure stpd stpd_name ports auto-edge [on | off] port_list
```

Description

Enables and disables auto-edge detection.

Syntax Description

<i>stpd_name</i>	Specifies an <i>STPD</i> name on the switch.
on	Enables auto-edge detection on the specified port.
off	Disables auto-edge detection on the specified port.
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

By default, auto-edge detection is on.

Usage Guidelines

None.

Example

The following example enables auto-edge detection on port 1:10 in *STP* domain r1:

```
configure stpd r1 ports auto-edge on 1:10
```

History

This command was first available in ExtremeXOS 15.7.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure stpd ports bpdu-restrict

```
configure {stpd} stpd_name ports bpdu-restrict [enable | disable]  
    port_list {recovery-timeout {seconds}}
```

Description

Configures BPDU Restrict.

Syntax Description

<i>stpd_name</i>	Specifies an <i>STPD</i> name on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.
bpdu-restrict	Disables port as soon as a BPDU is received.
recovery-timeout	Time after which the port will be re-enabled.
<i>seconds</i>	Specifies the time in seconds. The range is 60 to 600. The default is 300.

Default

The default is disabled.

Usage Guidelines

Before using this command, the port(s) should be configured for edge-safeguard.

Example

The following command enables bpdu-restrict on port 2 of STPD s1:

```
configure stpd s1 ports bpdu-restrict enable 2
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure stpd ports cost

```
configure stpd stpd_name ports cost [auto | cost] port_list
```


Description

Specifies the path cost of the port in the specified *STPD*.

Syntax Description

<i>stpd_name</i>	Specifies an STPD name on the switch.
auto	Specifies the switch to remove any user-defined port cost value(s) and use the appropriate default port cost value(s).
<i>cost</i>	Specifies a numerical port cost value. The range is 1 through 200,000,000.
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

The switch automatically assigns a default path cost based on the speed of the port, as follows:

- 10 Mbps port—the default cost is 2,000,000.
- 100 Mbps port—the default cost is 200,000.
- 1000 Mbps port—the default cost is 20,000.
- 10000 Mbps ports—the default cost is 2,000.

The default port cost for trunked ports is dynamically calculated based on the available bandwidth.

Usage Guidelines

If your STPD has the same name as another component, for example a *VLAN*, we recommend that you specify the identifying keyword as well as the name. If your STPD has a name unique only to that STPD, the keyword *stpd* is optional.

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 802.1D-2004 standard modified the default port path cost value to allow for higher link speeds. If you have a network with both 802.1D-2004 and 802.1D-1998 compliant bridges, a higher link speed can create a situation whereby an 802.1D-1998 compliant bridge could become the most favorable transit path and possibly cause the traffic to span more bridges. To prevent this situation, configure the port path cost to make links with the same speed use the same path host value. For example, if you have 100 Mbps links on all bridges, configure the port path cost for the 802.1D-2004 compliant bridges to 19 instead of using the default 200,000.



Note

You cannot configure the port path cost on 802.1D-1998 compliant bridges to 200,000 because the path cost range setting is 1 to 65,535.

The range for the cost parameter is 1 through 200,000,000. If you configure the port cost, a setting of 1 indicates the highest priority.

If you configured a port cost value and specify the auto option, the switch removes the user-defined port cost value and returns to the default, automatically assigned, port cost value.

The auto port cost of a trunk port is calculated based on number member ports in the trunk port. Link up and down of the member port does not affect the trunk port cost, thus it does not trigger topology change. Only adding or removing a member port to/from the trunk port causes auto trunk port cost to change. Also, by so configuring a static trunk port cost, the value is frozen regardless of the number of member ports in the trunk port.

ExtremeXOS 11.5 and Earlier

If you have switches running ExtremeXOS 11.5 and earlier, the default costs are different than switches running ExtremeXOS 11.6 and later.

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, as follows:

- 10 Mbps port—the default cost is 100.
- 100 Mbps port—the default cost is 19.
- 1000 Mbps port—the default cost is 4.
- 10000 Mbps ports—the default cost is 2.

Example

The following command configures a cost of 100 to slot 2, ports 1 through 5 in STPD s0:

```
configure stpd s0 ports cost 100 2:1-2:5
```

History

This command was first available in ExtremeXOS 10.1.

The auto option was added in ExtremeXOS 11.0.

The default costs were updated based on support for the 802.1D-2004 standard in ExtremeXOS 11.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure stpd ports edge-safeguard disable

```
configure {stpd} stpd_name ports edge-safeguard disable port_list {bpdu-  
restrict} {recovery-timeout {seconds}}
```

Description

Disables the edge safeguard loop prevention on the specified RSTP or *MSTP* edge port.

Syntax Description

<i>stpd_name</i>	Specifies an <i>STPD</i> name on the switch.
<i>port_list</i>	Specifies one or more edge ports.
bpdu-restrict	Disables port as soon as a BPDU is received.
recovery-timeout	Time after which the port will be re-enabled.
<i>seconds</i>	Specifies the time in seconds. The range is 60 to 600. The default is 300.

Default

By default, this feature is disabled.

Usage Guidelines

This command applies only to ports that have already been configured as edge ports.

Loop prevention and detection on an edge port configured for RSTP or MSTP is called edge safeguard. An edge port configured with edge safeguard immediately enters the forwarding state and transmits BPDUs.

If you disable this feature, the edge port enters the forwarding state but no longer transmits BPDUs unless a BPDU is received by that edge port. This is the default behavior.

Recovery time starts as soon as the port becomes disabled. If no recovery-timeout is specified, the port is permanently disabled.

BPDU restrict can be disabled using the `configure stpd stpd_name ports bpdu-restrict disableport_list` command.

If edge safeguard is disabled, BPDU restrict is also disabled.

To view the status of the edge safeguard feature use the `show {stpd} stpd_name ports {[detail | port_list {detail}]}` command. You can also use the `show stpd {stpd_name | detail}` command to display the STPD configuration on the switch, including the enable/disable state for edge safeguard.



Note

In MSTP, configuring edge safeguard at CIST will be inherited in all *MSTI*.

To enable or re-enable edge safeguard, use one of the following commands:

- `configure {stpd} stpd_name ports edge-safeguard enableport_list {bpdu-restrict} {recovery-timeout {seconds}}`
- `configure stpd stpd_name ports link-type [[auto | broadcast | point-to-point]port_list | edgeport_list {edge-safeguard [enable | disable] {bpdu-restrict} {recovery-timeoutseconds}}]`

Example

The following command disables edge safeguard on RSTP edge port 4 in STPD s1 on a stand-alone switch:

```
configure stpd s1 ports edge-safeguard disable 4
```

History

This command was first available in ExtremeXOS 11.4.

The BPDU Restrict function was added in ExtremeXOS 12.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure stpd ports edge-safeguard enable

```
configure {stpd} stpd_name ports edge-safeguard enable port_list {bpdu-restrict} {recovery-timeout {seconds}}
```

Description

Enables the edge safeguard loop prevention on the specified RSTP or MSTP edge port.

Syntax Description

<i>stpd_name</i>	Specifies an <u>STPD</u> name on the switch.
<i>port_list</i>	Specifies one or more edge ports.
bpdu-restrict	Disables port as soon as a BPDU is received.
recovery-timeout	Time after which the port will be re-enabled.
<i>seconds</i>	Specifies the time in seconds. The range is 60 to 600. The default is 300.

Default

By default, this feature is disabled.

Usage Guidelines

This command applies only to ports that have already been configured as edge ports.

Loop prevention and detection on an edge port configured for RSTP or MSTP is called edge safeguard. You configure edge safeguard on RSTP or MSTP edge ports to prevent accidental or deliberate misconfigurations (loops) resulting from connecting two edge ports together or by connecting a hub or

other non-*STP* switch to an edge port. Edge safeguard also limits the impact of broadcast storms that might occur on edge ports.

An edge port configured with edge safeguard immediately enters the forwarding state and transmits BPDUs. This advanced loop prevention mechanism improves network resiliency but does not interfere with the rapid convergence of edge ports.

Recovery time starts as soon as the port becomes disabled. If no recovery-timeout is specified, the port is permanently disabled.

BPDUs restrict can be disabled using the `configure {stpd} stpd_name ports bpdu-restrict [enable | disable]port_list {recovery-timeout {seconds}}` command and selecting disable.

If edge safeguard is disabled, BPDUs restrict is also disabled.

To view the status of the edge safeguard feature use the `show {stpd} stpd_name ports {[detail |port_list {detail}]}` command. You can also use the `show stpd {stpd_name | detail}` command to display the STPD configuration on the switch, including the enable/disable state for edge safeguard.



Note

In MSTP, configuring edge safeguard at CIST will be inherited in all *MSTI*.

To disable edge safeguard, use one of the following commands:

- `configure {stpd} stpd_name ports edge-safeguard disableport_list {bpdu-restrict} {recovery-timeout {seconds}}`
- `configure stpd stpd_name ports link-type [[auto | broadcast | point-to-point]port_list | edgeport_list {edge-safeguard [enable | disable] {bpdu-restrict} {recovery-timeoutseconds}}]`

Example

The following command enables edge safeguard on RSTP edge port 4 in STPD s1 on a stand-alone switch:

```
configure stpd s1 ports edge-safeguard enable 4
```

History

This command was first available in ExtremeXOS 11.4.

The BPDUs Restrict function was added in ExtremeXOS 12.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure stpd ports link-type

```
configure stpd stpd_name ports link-type [[auto | broadcast | point-to-point] port_list | edge port_list {edge-safeguard [enable | disable] {bpdu-restrict} {recovery-timeout seconds}}]
```

Description

Configures the ports in the specified *STPD* as auto, broadcast, edge, or point-to-point link types.

Syntax Description

<i>stpd_name</i>	Specifies an STPD name on the switch.
auto	Specifies the switch to automatically determine the port link type. An auto link behaves like a point-to-point link if the link is in full-duplex mode or if link aggregation is enabled on the port. Used for 802.1w configurations.
broadcast	Specifies a port attached to a LAN segment with more than two bridges. Used for 802.1D configurations. A port with broadcast link type cannot participate in rapid reconfiguration using RSTP or <i>MSTP</i> . By default, all STP.1D ports are broadcast links.
point-to-point	Specifies a port attached to a LAN segment with only two bridges. A port with point-to-point link type can participate in rapid reconfiguration. Used for 802.1w and MSTP configurations. By default, all 802.1w and MSTP ports are point-to-point link types.
<i>port_list</i>	Specifies one or more ports or slots and ports.
edge	Specifies a port that does not have a bridge attached. An edge port is placed and held in the <i>STP</i> forwarding state unless a BPDU is received by the port. Used for 802.1w and MSTP configurations.
edge-safeguard	Specifies that the edge port be configured with edge safeguard, a loop prevention and detection mechanism. Used for 802.1w and MSTP configurations.
enable	Specifies that edge safeguard be enabled on the edge port(s).
disable	Specifies that edge safeguard be disabled on the edge port(s).
bpdu-restrict	Disables port as soon as a BPDU is received.
recovery-timeout	Time after which the port will be re-enabled.
<i>seconds</i>	Specifies the time in seconds. The range is 60 to 600. The default is 300.

Default

STP.1D ports are broadcast link types 802.1w and MSTP ports are auto link types.

Usage Guidelines

If your STPD has the same name as another component, for example a VLAN, we recommend that you specify the identifying keyword as well as the name. If your STPD has a name unique only to that STPD, the keyword `stpd` is optional.

The default, broadcast links, supports legacy STP (802.1D) configurations. If the switch operates in 802.1D mode, any configured port link type will behave the same as the broadcast link type.

RSTP rapidly moves the designated ports of a point-to-point link type into the forwarding state. This behavior is supported by RSTP and MSTP only.

In an MSTP environment, configure the same link types for the CIST and all MSTIs.

Auto Link Type

An auto link behaves like a point-to-point link if the link is in full duplex mode or if link aggregation is enabled on the port; otherwise, an auto link behaves like a broadcast link. If a non-STP switch exists between several switches operating in 802.1w mode with auto links, the non-STP switch may negotiate full-duplex even though the broadcast domain extends over several STP devices.

Edge Link Type

RSTP does not send any BPDUs from an edge port nor does it generate topology change events when an edge port changes its state.

If you configure a port to be an edge port, the port immediately enters the forwarding state. Edge ports remain in the forwarding state unless the port receives a BPDU. In that case, edge ports enter the blocking state. The edge port remains in the blocking state until it stops receiving BPDUs and the message age timer expires.

Edge Safeguard

Loop prevention and detection on an edge port configured for RSTP or MSTP is called edge safeguard. You configure edge safeguard on RSTP or MSTP edge ports to prevent accidental or deliberate misconfigurations (loops) resulting from connecting two edge ports together or by connecting a hub or other non-STP switch to an edge port. Edge safeguard also limits the impact of broadcast storms that might occur on edge ports.

An edge port configured with edge safeguard immediately enters the forwarding state and transmits BPDUs. This advanced loop prevention mechanism improves network resiliency but does not interfere with the rapid convergence of edge ports.

Recovery time starts as soon as the port becomes disabled. If no recovery-timeout is specified, the port is permanently disabled.

BPDU restrict can be disabled using the `configure stpd stpd_name ports bpdu-restrict disableport_list` command.

If edge safeguard is disabled, BPDU restrict is also disabled.

To configure a port as an edge port and enable edge safeguard on that port, use the `configure stpd stpd_name ports link-type edgeport_list edge-safeguard` command and specify `enable`.

To disable edge safeguard on the edge port, use the `configure stpd stpd_name ports link-type edgeport_list edge-safeguard` command and specify `disable`.

Two other commands are also available to enable and disable edge safeguard:

```
configure stpd ports edge-safeguard enable
configure stpd ports edge-safeguard disable
```

In MSTP, configuring edge safeguard at CIST will be inherited in all *MSTI*.

Example

The following command configures slot 2, ports 1 through 4 to be point-to-point links in STPD s1:

```
configure stpd s1 ports link-type point-to-point 2:1-2:4
```

The following command enables edge safeguard on the RSTP edge port on slot 2, port 3 in STPD s1 configured for RSTP:

```
configure stpd s1 ports link-type edge 2:3 edge-safeguard enable
```

History

This command was first available in ExtremeXOS 10.1.

The BPDU Restrict function was added in ExtremeXOS 12.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure stpd ports loop-protect

```
configure stpd stpd_name ports loop-protect [on | off] port_list
```

Description

Enables and disables loop protect on a port.

Syntax Description

<i>stpd_name</i>	Specifies an <i>STPD</i> name on the switch.
on	Enables loop protect on the specified port.

off	Disables loop protect on the specified port.
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

By default, loop protect is **off**.

Usage Guidelines

Loop protect prevents loops due to misconfiguration or one-way communication failures.

Example

The following example enables loop protect on port 1:10 in the *STP* domain r1:

```
configure stpd r1 ports loop-protect on 1:10
```

History

This command was first available in ExtremeXOS 15.7.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure stpd ports loop-protect partner

```
configure stpd stpd_name ports loop-protect partner [capable |  
incapable] port_list
```

Description

Configures whether the link partner is capable of the loop protect feature.

Syntax Description

<i>stpd_name</i>	Specifies an <i>STPD</i> name on the switch.
capable	The link partner supports the loop protect feature.
incapable	The link partner does not support the loop protect feature.
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

By default, this command is set to **incapable**.

Usage Guidelines

Ports work in two loop protect operational modes:

- If the port is set to **capable**, the port works in full mode.
- If the port is set to **incapable**, the port works limited mode.

In full mode, when RSTP/*MSTP* BPDUs are received on a point-to-point link and the port is designated, a loop protect timer is set to three times the hello time. When this timer expires, the port is moved to the blocking state. Limited mode adds the requirement that the **flags** field in the BPDU indicates a root role.

Example

The following example configures loop protect partner capability to "capable" for port 1:10 in the *STP* domain r1:

```
configure stpd r1 ports loop-protect partner capable 1:10
```

History

This command was first available in ExtremeXOS 15.7.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure stpd ports mode

```
configure stpd stpd_name ports mode [dot1d | emistp | pvst-plus]  
port_list
```

Description

Configures the encapsulation mode for the specified port list.

Syntax Description

<i>stpd_name</i>	Specifies an <i>STPD</i> name on the switch.
dot1d	Specifies the <i>STP</i> encapsulation mode of operation to be 802.1d.
emistp	Specifies the STP encapsulation mode of operation to be EMISTP.
pvst-plus	Specifies the STP encapsulation mode of operation to be PVST+.
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

Ports in the default STPD (s0) and user-created STPDs are dot1d mode.

Usage Guidelines

If your STPD has the same name as another component, for example a VLAN, we recommend that you specify the identifying keyword as well as the name. If your STPD has a name unique only to that STPD, the keyword `stpd` is optional.

MSTP STPDs use 802.1D BPDU encapsulation mode by default. To ensure correct operation of your MSTP STPDs, do not configure EMISTP or PVST+ encapsulation mode for MSTP STPDs.

You can specify the following STP encapsulation modes:

- `dot1d`—This mode is reserved for backward compatibility with previous STP versions. BPDUs are sent untagged in 802.1D mode. Because of this, any given physical interface can have only one STPD running in 802.1D mode.

This encapsulation mode supports the following STPD modes of operation: 802.1D, 802.1w, and MSTP.

- `emistp`—This mode sends BPDUs with an 802.1Q tag having an STPD ID in the VLAN ID field.

This encapsulation mode supports the following STPD modes of operation: 802.1D and 802.1w.

- `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.

This encapsulation mode supports the following STPD modes of operation: 802.1D and 802.1w.

Example

The following command configures STPD `s1` with PVST+ packet formatting for slot 2, port 1:

```
configure stpd s1 ports mode pvst-plus 2:1
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure stpd ports port-priority

```
configure stpd stpd_name ports port-priority priority port_list
```

Description

Specifies the port priority of the port in the specified STPD.

Syntax Description

<i>stpd_name</i>	Specifies an STPD name on the switch.
<i>priority</i>	Specifies a numerical port priority value. The range is 0 through 240 and is subject to the multiple of 16 restriction.
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

The default is 128.

Usage Guidelines

If your STPD has the same name as another component, for example a [VLAN](#), we recommend that you specify the identifying keyword as well as the name. If your STPD has a name unique only to that STPD, the keyword `stpd` is optional.

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.

To preserve backward compatibility and to use ExtremeXOS 11.5 or earlier configurations, the existing `configure stpd ports priority` command is available in ExtremeXOS 11.6. If you have an ExtremeXOS 11.5 or earlier configuration, the switch interprets the port priority based on the 802.1D-1998 standard. If the switch reads a value that is not supported in ExtremeXOS 11.6, the switch rejects the entry. For example, if the switch reads the `configure stpd ports priority 16` command from an ExtremeXOS 11.5 or earlier configuration, (which is equivalent to the command `configure stpd ports priority 8` entered through CLI), the switch saves the value in the new ExtremeXOS 11.6 configuration as `configure stpd ports port-priority 128`.

A setting of 0 indicates the highest priority.

The range for the priority parameter is 0 through 240 and is subject to the multiple of 16 restriction.

Example

The following command assigns a priority of 32 to slot 2, ports 1 through 5 in STPD s0:

```
configure stpd s0 ports port-priority 32 2:1-2:5
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure stpd ports priority

configure stpd *stpd_name* **ports priority** *priority* *port_list*

Description

Specifies the port priority of the port in the specified *STPD*.

Syntax Description

<i>stpd_name</i>	Specifies an STPD name on the switch.
<i>priority</i>	Specifies a numerical port priority value. The range is 0 through 31 for <i>STP</i> and 0 through 15 for <i>MSTP</i> and RSTP.
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

The default is 128.

Usage Guidelines

If your STPD has the same name as another component, for example a *VLAN*, we recommend that you specify the identifying keyword as well as the name. If your STPD has a name unique only to that STPD, the keyword stpd is optional.

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.

To preserve backward compatibility and to use ExtremeXOS 11.5 or earlier configurations, the existing `configure stpd ports priority` command is available in ExtremeXOS 11.6. If you have an ExtremeXOS 11.5 or earlier configuration, the switch interprets the port priority based on the 802.1D-1998 standard. If the switch reads a value that is not supported in ExtremeXOS 11.6, the switch rejects the entry.

A setting of 0 indicates the highest priority.

The range for the priority parameter is 0 through 31 for STP and 0 through 15 for MSTP and RSTP.

ExtremeXOS 11.6 introduces support for a new ports priority command: `configure stpd ports port-priority`. When you save the port priority value in an ExtremeXOS 11.6 configuration, the switch saves it as the new command `configure stpd ports port-priority` with the corresponding change in priority values. The priority range of this command is 0 through 240 and is subject to the multiple of 16 restriction. For more information see `configure stpd ports port-priority`.

ExtremeXOS 11.5 and Earlier

If you have switches running ExtremeXOS 11.5 and earlier, the default value for the priority range are different than switches running ExtremeXOS 11.6.

The range for the priority parameter is 0 through 31.

The default is 16.

Example

The following command assigns a priority of 1 to slot 2, ports 1 through 5 in STPD s0:

```
configure stpd s0 ports priority 1 2:1-2:5
```

History

This command was first available in ExtremeXOS 10.1.

The priority range and behavior was updated based on support for the 802.1D-2004 standard in ExtremeXOS 11.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure stpd ports reflection-bpdu

```
configure stpd stpd_name ports reflection-bpdu [on | off] port_list
```

Description

Turns on/off reflection Bridge Protocol Data Unit (BPDU) behavior.

Syntax Description

stpd	Spanning Tree Protocol (STP) domain.
<i>stpd_name</i>	Specifies the STP domain name
ports	Ports in this STP domain to configure.
reflection-bpdu	Copy contents (bridge ID, root ID, etc.) of received RSTP/MSTP proposal BPDU in transmitted agreement BPDU. Default is on.
on	Use received bridge ID, etc. in agreement RSTP/MSTP BPDU (not necessary for OUIs 00:01:F4, 00:11:88, 00:1F:45, 20:B3:99).
off	Use local bridge ID, etc. in transmitted agreement RSTP/MSTP BPDU for compatibility with EOS switches with unknown OUIs.
<i>port_list</i>	Specifies the ports in this STP domain to configure.

Default

Reflection BPDU behavior is on by default.

Usage Guidelines

For Rapid Spanning Tree Protocol (RSTP) proposal handshake to work with CISCO switches, the switch that receives the proposal BPDU reflects back the same BPDU (all the contents) with an agreement flag set. This ensures that the other port is acknowledging the proposal that the switch has send out, so the acknowledgment BPDU contains the same contents of the other switch's proposal BPDU with the agreement bit set, instead of the proposal bit.

However, this behavior when used with EOS upstream bridges receiving the agreement BPDU (whose MAC OUI is different than 00:01:F4, 00:11:88, 00:1F:45, 20:B3:99) causes the switch to believe it is being sent its own BPDU, thus causing a multisource event during a topology change. This command allows you turn off the BPDU reflection behavior to avoid this problem.

Example

To enable reflection BPDU on domain "s1" on port 7:

```
configure s1 ports reflection-bpdu on 7
```

To disable reflection BPDU on domain "s1" on port 7:

```
configure s1 ports reflection-bpdu off 7
```

History

This command was first available in ExtremeXOS 22.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure stpd ports restricted-role disable

```
configure stpd stpd_name ports restricted-role disable port_list
```

Description

Disables restricted role on the specified port inside the core network.

Syntax Description

<i>stpd_name</i>	Specifies an <i>STPD</i> name on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

N/A.

Usage Guidelines

The restricted role is disabled by default. If set, it can cause a lack of spanning tree connectivity. A network administrator enables the restricted role to prevent bridges external to a core region of the network from influencing the spanning tree active topology, possibly because those bridges are not under the full control of the administrator.



Note
Disabling Restricted Role at CIST is inherited by all *MSTI*.

Example

The following command disables restricted role for s1 on port 6:3:

```
configure stpd s1 ports restricted-role disable 6:3
```

History

This command was first available in ExtremeXOS 12.1.

This command was added to RSTP in ExtremeXOS 11.6 and 12.0.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure stpd ports restricted-role enable

```
configure stpd stpd_name ports restricted-role enable port_list
```

Description

Enables restricted role on the specified port inside the core network.

Syntax Description

<i>stpd_name</i>	Specifies an <i>STPD</i> name on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

N/A.

Usage Guidelines

Enabling restricted role causes the port not to be selected as a root port even if it has the best spanning tree priority vector. Such a port is selected as an alternate port after the root port has been selected.

The restricted role is disabled by default. If set, it can cause a lack of spanning tree connectivity. A network administrator enables the restricted role to prevent bridges external to a core region of the network from influencing the spanning tree active topology, possibly because those bridges are not under the full control of the administrator.



Note

Restricted role should not be enabled with edge mode.

Enabling Restricted Role at CIST is inherited by all MSTI.

Example

The following command enables restricted role on port 6:3:

```
configure stpd s1 ports restricted-role enable 6:3
```

History

This command was first available in ExtremeXOS 12.1.

This command was added to RSTP in ExtremeXOS 11.6 and 12.0.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure stpd ports restricted-tcn

```
configure stpd stpd_name ports restricted-tcn [on | off] port_list
```

Description

Restricts the propagation of Topology Change Notification (TCN) BPDUs on the specified port.

Syntax Description

<i>stpd_name</i>	Specifies an <u>STPD</u> name on the switch.
on	Does not propagate received TCN BPDUs and topology changes to other ports.
off	Allows the propagation of received TCN BPDUs and topology changes to other ports.
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

The default value is **off**.

Usage Guidelines

Set **restricted-tcn** to **on** to prevent unnecessary address flushing caused by persistent TCNs. Restricting TCNs is a useful when it is not possible to remove the source of the TCNs.

Example

The following example disables the propagation of TCNs in port 1:10 for *STP* domain r1:

```
configure stpd r1 ports restricted-tcn on 1:10
```

History

This command was first available in ExtremeXOS 15.7.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure stpd priority

```
configure stpd stpd_name priority priority
```

Description

Specifies the bridge priority of the *STPD*.

Syntax Description

<i>stpd_name</i>	Specifies an STPD name on the switch.
<i>priority</i>	<p>Specifies the bridge priority of the STPD. The range is 0 through 61,440.</p> <ul style="list-style-type: none">• If the bridge priority mode is configured as dot1d and the protocol mode is configured as dot1w, then value can be configured in increments of 1.• If the bridge priority mode is configured as dot1t and the protocol mode is configured as dot1w, then priority value can be configured in increments of 4,096.

Default

The default priority is 32,768.

Usage Guidelines

If your STPD has the same name as another component, for example a VLAN, we recommend that you specify the identifying keyword as well as the name. If your STPD has a name unique only to that STPD, the keyword **stpd** is optional.

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 priority parameter is 0 through 61,440. If the bridge priority mode is configured as dot1d and the protocol mode is configured as dot1w, then value can be configured in increments of 1. If the bridge priority mode is configured as dot1t and the protocol mode is configured as dot1w, then priority value can be configured in increments of 4,096. A setting of 0 indicates the highest priority.

If you have an ExtremeXOS 11.5 or earlier configuration that contains an STP or RSTP bridge priority that is not a multiple of 4,096, the switch rejects the entry and the bridge priority returns to the default value. The MSTP implementation already uses multiples of 4,096 to determine the bridge priority.

For example, to lower the numerical value of the priority (which gives the priority a higher precedence), you subtract 4,096 from the default priority: $32,768 - 4,096 = 28,672$. If you modify the priority by a value other than 4,096, the switch rejects the entry.

ExtremeXOS 11.5 and Earlier

If you have switches running ExtremeXOS 11.5 and earlier, the priority range is different than switches running ExtremeXOS 11.6 and later.

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:

```
configure stpd stpd1 priority 16384
```

History

This command was first available in ExtremeXOS 10.1.

The priority range and behavior was updated based on support for the 802.1D-2004 standard in ExtremeXOS 11.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure stpd priority-mode

```
configure stpd stpd_name priority-mode [dot1d | dot1t]
```

Description

Sets *STP* bridge priority values.

Syntax Description

stp	STP domain/STP global configuration.
<i>stp_name</i>	STP domain name on the switch.
priority-mode	Control allowable bridge priority values.
dot1d	Allow any bridge priority value. Valid values are 0–65,535 (in increments of 1), with 0 indicating high priority and 65,535 low priority.
dot1t	Allow bridge priority in steps of 4,096. This option is the default bridge priority mode. Valid values are 0–61,440 (in increments of 4,096), with 0 indicating high priority and 61,440 low priority. Values are automatically rounded up or down depending on the dot1t value to which the entered value is closest.

Default

dot1t option is configured by default for operation mode dot1w and *MSTP*.

Example

The following example configures the priority-mode as dot1d:

```
configure stp s1 priority-mode dot1d
```

History

This command was first available in ExtremeXOS 22.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure stpd tag

```
configure stpd stpd_name tag stpd_tag
```

Description

Assigns an StpdID to an *STPD*.

Syntax Description

<i>stpd_name</i>	Specifies an STPD name on the switch.
<i>stpd_tag</i>	Specifies the <u>VLAN</u> ID of the carrier VLAN that is owned by the STPD.

Default

N/A.

Usage Guidelines

If your STPD has the same name as another component, for example a VLAN, we recommend that you specify the identifying keyword as well as the name. If your STPD has a name unique only to that STPD, the keyword `stpd` is optional.

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 STPD ID is used to identify each STP domain. You assign the `StpdID` when configuring the domain. An STPD ID must be identical to the VLAN ID of the carrier VLAN in that STP domain, and that VLAN cannot belong to another STPD. Unless all ports are running in 802.1D mode, an STPD with ports running in either EMISTP mode or PVST+ mode must be configured with an STPD ID.

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 `configure vlan` commands.

MSTP Only

MSTP uses two different methods to identify the STPDs that are part of the MSTP network. An instance ID of 0 identifies the CIST. The switch assigns this ID automatically when you configure the CIST STPD. To configure the CIST STPD, use the `configure stpd stpd_name mode [dot1d | dot1w | mstp [cist | mstiinstance]]` command.

An MSTI identifier (MSTI ID) identifies each STP domain that is part of an MSTP region. You assign the MSTI ID when configuring the STPD that participates in the MSTP region. Each STPD that participates in a particular MSTP region must have the same MSTI ID. To configure the MSTI ID, use the `configure stpd stpd_name mode [dot1d | dot1w | mstp [cist | mstiinstance]]` command.

Example

The following example assigns an `StpdID` to the `purple_st` STPD:

```
configure stpd purple_st tag 200
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure stpd trap new-root

```
configure stpd stpd_name trap new-root [on | off]
```

Description

Enables and disables the new-root trap.

Syntax Description

<i>stpd_name</i>	Specifies an <i>STPD</i> name on the switch.
on	Enables the new-root trap.
off	Disables the new-root trap.

Default

By default, the trap is enabled (**on**).

Usage Guidelines

The new-root trap is sent when the new root bridge is elected.

Example

The following example disables the new-root trap for the *STP* domain r1.

```
configure stpd r1 trap new-root off
```

History

This command was first available in ExtremeXOS 15.7.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure stpd trap topology-change

```
configure stpd stpd_name trap topology-change {edge-ports} [on | off]
```

Description

Enables and disables the topology change trap for all ports or edge ports only.

Syntax Description

<i>stpd_name</i>	Specifies an <i>STPD</i> name on the switch.
edge-ports	Specifies that topology change traps will be sent only for edge ports.
on	Enables the topology change trap.
off	Disables the topology change trap.

Default

By default, the topology change trap is disabled (**off**) for all ports.

Usage Guidelines

You cannot enable the topology change trap for edge ports if you have disabled the topology change trap for all ports.

Example

The following example disables the topology change trap for edge ports only in the *STP* domain r1.

```
configure stpd r1 trap topology-change edge-ports off
```

History

This command was first available in ExtremeXOS 15.7.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure stpd tx-hold-count

```
configure stpd stpd_name tx-hold-count tx_hold_count
```

Description

Configures the maximum BPDUs transmitted per second.

Syntax Description

<i>stpd_name</i>	Specifies an <i>STPD</i> name on the switch.
<i>tx_hold_count</i>	Specifies the maximum number of BPDUs transmitted per second. The valid range is 1-10.

Default

By default, the maximum number of BPDUs transmitted per second is 6.

Usage Guidelines

The transmit hold count is used by the port transmit state machine to limit BPDU transmission rate.

Example

The following example configures the transmit hold count for *STP* domain r1 to five BPDUs per second:

```
configure stpd r1 tx-hold-count 5
```

History

This command was first available in ExtremeXOS 15.7.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure switch integrity-check image

```
configure switch integrity-check image [on | off]
```

Description

Enables or disables the NOS image integrity check feature.

Syntax Description

integrity-check	Specifies configuring integrity check.
image	Specifies checking the loaded NOS image.
on	Enables image integrity check feature. You are informed if the integrity check fails.
off	Disables image integrity check feature (default).

Default

By default, ExtremeXOS image integrity check is disabled.

Usage Guidelines

If the image integrity check is enabled, during bootup, the system checks the integrity of the NOS image, and notifies you if it has been compromised or not (an error message is logged).

To view the status and configuration of the image integrity check, use the `show switch management` command.

Example

The following example enables the NOS image integrity check:

```
# configure switch integrity-check image on
```

History

This command was first available in ExtremeXOS 31.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure sys-health-check all level

```
configure sys-health-check all level [normal | strict]
```

Description

Configures how the ExtremeXOS software handles faults for the switch.

Syntax Description

normal	Upon a fault detection, the switch only sends a message to the syslog. This is the default setting.
strict	Upon a fault detection, the switch takes the action configured by the <code>configure sys-recovery-level slot</code> or the command.

Default

The default setting is normal.

Usage Guidelines

Use this command in conjunction with the `configure sys-recovery-level switch [none | reset | shutdown]` command to implement your network’s fault handling strategy.

ExtremeXOS 11.5 enhances the number of switch-fabric tests completed and monitored by the polling module of the system health checker. Additionally with ExtremeXOS 11.5, you can now configure how ExtremeXOS handles a detected fault based on the configuration of the `configure sys-recovery-level slot [all | slot_number] [none | reset | shutdown]` or the `configure sys-recovery-level switch [none | reset | shutdown]` command.

If you configure the strict parameter, the switch takes the action configured by the `configure sys-recovery-level slot` or the `configure sys-recovery-level switch` command, which can include logging only or restarting, rebooting, or shutting down the suspect device.

To maintain a smooth upgrade for devices running ExtremeXOS 11.4 and earlier, the switch-fabric tests introduced in ExtremeXOS 11.5 are set to only log error messages ('normal mode') by default. However, we recommend that you configure 'strict mode' so the system can attempt to recover by utilizing the action configured in the `configure sys-recovery-level slot` or the `configure sys-recovery-level switch` command (which by default is reset).

Depending on your switch configuration, the following table shows how ExtremeSwitching series switches behave when the ExtremeXOS software detects a fault:

Table 19: System Behavior for ExtremeSwitching Series Switches

Fault Handling Configuration	Hardware Recovery Configuration	Behavior
<code>configure sys-health-check all level normal</code>	<code>configure sys-recovery-level switch none</code>	The switch sends messages to the syslog.
Same as above.	<code>configure sys-recovery-level switch reset</code>	Same as above.
Same as above.	<code>configure sys-recovery-level switch shutdown</code>	Same as above.
<code>configure sys-health-check all level strict</code>	<code>configure sys-recovery-level switch none</code>	Same as above.
Same as above.	<code>configure sys-recovery-level switch reset</code>	ExtremeXOS reboots the affected switch.
Same as above.	<code>configure sys-recovery-level switch shutdown</code>	ExtremeXOS shuts down the affected switch.

Displaying the System Health Check Setting

To display the system health check setting, including polling and how ExtremeXOS handles faults on the switch, use the following command:

```
show switch
```

The system health check setting, displayed as SysHealth check, shows the polling setting and how ExtremeXOS handles faults. The polling setting appears as Enabled, and the fault handling setting appears in parenthesis next to the polling setting. In the following truncated output, the system health check setting appears as SysHealth check: Enabled (Normal):

```
SysName:          TechPubs Lab
SysName:          BD-8810Rack3
SysLocation:
SysContact:       support@extremenetworks.com, +1 888 257 3000
System MAC:       00:04:96:1F:A2:60
SysHealth check:  Enabled (Normal)
Recovery Mode:    None
System Watchdog:  Enabled
```

If you use the strict parameter, which configures the switch to take the action configured by the `configure sys-recovery-level slot` or the `configure sys-recovery-level switch` command, (Strict) would appear next to Enabled.

Example

The following command configures the switch to forward faults to be handled by the level set by the `configure sys-recovery-level switch` command:

```
# configure sys-health-check all level strict
```

History

This command was first available in ExtremeXOS 11.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure syslog add

```
configure syslog add [ipaddress {udp-port {udp_port}} | ipPort |  
                    ipaddress tls_port {tls_port}] {vr vr_name} [local0...local7]
```

Description

Configures the remote Syslog server host address, and filters messages to be sent to the remote Syslog target.

Syntax Description

<i>ipaddress</i>	Specifies the remote Syslog server IP address.
<i>ipPort</i>	Specifies the UDP port number for the Syslog target.
tls_port	Specifies remote Syslog server Transport Layer Security (TLS) for connection type.
<i>tls_port</i>	TLS port number (default is 6514).
<i>vr_name</i>	Specifies the virtual router that can reach the server IP address. Note: User-created VRs are supported only on the platforms listed for this feature in the Switch Engine 32.3 Feature License Requirements document.
local0 ... local7	Specifies the local Syslog facility.

Default

If a virtual router is not specified, VR-Mgmt is used. If UDP port is not specified, 514 is used. If TLS port is not specified, 6514 is used.

Usage Guidelines

Options for configuring the remote Syslog server include:

- `ipaddress`—The IP address of the remote Syslog server host
- `ipPort`—The UDP port
- `vr_name`—The virtual router that can reach the Syslog host
- `local0-local7`—The Syslog facility level for local use

The switch log overwrites existing log messages in a wrap-around memory buffer, which may cause you to lose valuable information once the buffer becomes full. The remote Syslog server does not overwrite log information, and can store messages in non-volatile files (disks, for example).

The `enable syslog` command must be issued in order for messages to be sent to the remote Syslog server(s). Syslog is disabled by default. A total of four Syslog servers can be configured at one time.

When a Syslog server is added, it is associated with the filter `DefaultFilter`. Use the `configure log target filter` command to associate a different filter.

The Syslog facility level is defined as `local0 - local7`. The facility level is used to group Syslog data.

Example

The following example adds the remote Syslog server with an IP address of 10.0.0.1:

```
configure syslog add 10.0.0.1 local1
```

The following example adds the remote Syslog server with an IP address of 2001:11::123:

```
configure syslog add 2001:11::123 local1
```

History

This command was first available in ExtremeXOS 10.1.

The `ipPort` parameter was first available in ExtremeXOS 11.0.

The `udp-port` parameter and support for the *EMS (Event Management System)* to send log messages to Syslog servers having IPv6 address was added in ExtremeXOS 21.1.

Transport Layer Security (TLS) option added in ExtremeXOS 22.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure syslog tls cipher

```
configure syslog tls cipher [[cipher | all] on | cipher off]
```

Description

Turns on/off ciphers for Syslog Transport Layer Security (TLS) sessions.

Syntax Description

syslog	Specifies configuring the remote Syslog target.
tls	Transport Layer Security (TLS) protocol.
cipher	Specifies configuring the algorithm to use for encrypting Syslog TLS sessions.
<i>cipher</i>	Specifies the cipher name to enable or disable.
all	Specifies all ciphers for enabling.
on	Enable selected cipher. Default is that all ciphers are on.
off	Disables selected cipher.

Default

By default, all ciphers are enabled.

Usage Guidelines

A minimum of one cipher must be enabled.

The following is the list of available ciphers:

- aes128-sha
- aes128-sha256
- aes256-sha256
- dhe-rsa-aes128-sha256
- dhe-rsa-aes256-sha256

To view which ciphers are enabled and disabled, use the command [show log configuration](#) on page 2828.

Example

The following example enables all ciphers for Syslog TLS sessions:

```
configure syslog tls cipher all on
```

The following example disables the aes128-sha cipher for Syslog TLS sessions:

```
configure syslog tls cipher aes128-sha off
```

History

This command was first available in ExtremeXOS 22.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure syslog tls ocp

```
configure syslog tls ocp [on | off]
```

Description

Enables or disables Online Certificate Status Protocol (OCSP) check for Transport Layer Security (TLS) connections to remote Syslog servers.

Syntax Description

syslog	Specifies configuring the remote Syslog target.
tls	Specifies configuring TLS.
ocsp	Specifies configuring OCSP for real-time certificate revocation status checking.
on	Enables OCSP (default).
off	Disables OCSP.

Default

By default, OCSP is enabled.

Usage Guidelines

While you can disable OCSP, it is not recommended because no certificate revocation status check is performed.

Example

The following example enables OCSP check for TLS connections to remote Syslog servers.

```
# configure syslog tls ocp on
```

History

This command was first available in ExtremeXOS 30.7.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure syslog tls ocp nonce

```
configure syslog tls ocp nonce [on | off]
```

Description

Enables or disables Online Certificate Status Protocol (OCSP) nonce for Transport Layer Security (TLS) connections to remote Syslog servers.

Syntax Description

syslog	Specifies configuring the remote Syslog target.
tls	Specifies configuring TLS.
ocsp	Specifies configuring OCSP for real-time certificate revocation status checking.
nonce	Specifies to cryptographically bind an OCSP request and an OCSP response with the extension <code>id-pkix-ocsp-nonce</code> to prevent replay attacks.
on	Specifies to include the <code>id-pkix-ocsp-nonce</code> extension in the OCSP request and response.
off	Specifies to exclude the extension (default).

Default

Off.

Usage Guidelines

Example

The following example configures nonce:

```
# configure syslog tls ocsp nonce on
```

History

This command was first available in ExtremeXOS 32.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure syslog tls ocsp override

```
configure syslog tls tls override [url | none]
```

Description

This command configures one HTTP Online Certificate Status Protocol (OCSP) override URL for Transport Layer Security (TLS) connections to a remote Syslog server.

Syntax Description

syslog	Specifies configuring the remote Syslog target.
tls	Specifies Transport Layer Security (TLS).
ocsp	Specifies the OCSP attribute.
override	Specifies to override the OCSP server in the AuthorityInformationAccess section of a syslog server's certificate.
<i>url</i>	Specifies the URL of the OCSP override server. Default port is 80.
none	Specifies to remove the OCSP override URL configuration (default).

Default

None.

Usage Guidelines

Only HTTP is supported with either FQDN or IP.

Example

The following example configures an override URL of `http://syslogocsp:2022`:

```
# configure radius tls ocsp override http://syslogocsp:2022
```

History

This command was first available in ExtremeXOS 32.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure syslog tls ocsp signer

```
configure syslog tls ocsp signer ocsp-nocheck [on | off]
```

Description

Enables or disables Online Certificate Status Protocol (OCSP) signer's `ocsp-nocheck` for Transport Layer Security (TLS) connections to remote Syslog servers.

Syntax Description

syslog	Specifies configuring the remote Syslog target.
tls	Specifies Transport Layer Security (TLS).

ocsp	Specifies configuring OCSP for real-time certificate revocation status checking.
signer	Specifies the OCSP signer that signs the OCSP response.
ocsp-nocheck	Specifies the extension <code>id-pkix-ocsp-nocheck</code> . If present in the OCSP signer's certificate, then it is trusted for its lifetime.
on	Specifies to override the <code>id-pkix-ocsp-nocheck</code> extension in the OCSP signer's certificate and forces the extension as if it is present.
off	Specifies to behave per the extension's presence in the OCSP signer's certificate. If not present and the OCSP signer is not root CA, then the whole OCSP will fail (default).

Default

Off.

Usage Guidelines

Example

The following example enables OCSP signer's nocheck for TLS connections to a remote Syslog server.

```
# configure syslog tls ocsp signer ocsp-nocheck on
```

History

This command was first available in ExtremeXOS 32.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure syslog tls tcp-user-timeout

```
configure syslog tls tcp-user-timeout [seconds | default]
```

Description

Specifies the maximum time that transmitted data may remain unacknowledged before TCP closes the connection to avoid loss of logging to TLS Syslog server.

Syntax Description

tls	Specifies Transport Layer Security protocol.
tcp-user-timeout	Specifies the maximum time that transmitted data may remain unacknowledged before TCP closes the connection.

<i>seconds</i>	Timeout period in seconds. Range = 20–900.
default	Specifies not using value from tcp-user-timeout option; use the system default.

Default

The default is to use Linux default—tcp-user-timeout is not enabled.

Usage Guidelines

For Linux, by default, it takes about 15 minutes for kernel to end a TCP connection when transmitted data remains unacknowledged. This results in a potential loss of logs to TLS Syslog server during the 15 minutes window due to link down. This command allows you to reduce this window.

Example

The following example sets the TCP user timeout value to 30 seconds:

```
configure syslog tls tcp-user-timeout 30
```

The following example turns off using the TCP user timeout value and accepts system default:

```
configure syslog tls tcp-user-timeout default
```

History

This command was first available in ExtremeXOS 22.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure syslog delete

```
configure syslog delete [ ipaddress {udp-port {udp_port}} | ipPort |
    ipaddress tls_port {tls_port}] {vr vr_name} [local0...local17 ] ] | all
    {local0...local17} {vr vr_name} ]

configure syslog delete host name/ip {: udp-port} [local0...local17]
```

Description

Deletes a remote Syslog server address.

Syntax Description

<i>ipaddress</i>	Specifies the remote Syslog server IP address.
<i>ipPort</i>	Specifies the UDP port number for the Syslog target.

tls_port	Specifies remote Syslog server Transport Layer Security (TLS) for connection type.
<i>tls_port</i>	TLS port number.
<i>vr_name</i>	Specifies the virtual router that can reach the server IP address. Note: User-created VRs are supported only on the platforms listed for this feature in the Switch Engine 32.3 Feature License Requirements document.
local10 ... local17	Specifies the local Syslog facility.
all	Specifies all remote Syslog servers.

Default

If a virtual router is not specified, *VR-Mgmt* is used.

If a UDP port number is not specified, 514 is used.

If a TLS port number is not specified, 6514 is used.

Usage Guidelines

This command is used to delete a remote Syslog server target.

Example

The following example deletes the remote Syslog server with an IP address of 10.0.0.1:

```
configure syslog delete 10.0.0.1 local1
```

The following example deletes the remote Syslog server with an IP address of 2001:11::123 :

```
configure syslog delete 2001:11::123 local1
```

History

This command was first available in ExtremeXOS 10.1.

The **ipPort** parameter was first available in ExtremeXOS 11.0.

The **udp-port** parameter and support for the EMS to send log messages to Syslog servers having IPv6 address was added in ExtremeXOS 21.1.

Transport Layer Security (TLS) option added in ExtremeXOS 22.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure syslog reference-identifier

```
configure syslog [all | ipaddress {tls-port tls_port}] {vr vr_name}  
    {local} reference-identifier reference_identifier
```

Description

Specifies the remote Syslog server certificate reference identifier.

Syntax Description

all	All specified targets.
<i>ipaddress</i>	Specifies the remote Syslog server IPv4 or IPv6 address.
tls-port	Specifies using a remote Syslog server Transport Layer Security (TLS) port.
<i>tls_port</i>	Specifies the remote Syslog server Transport Layer Security (TLS) port (default is 6514).
vr	Specifies a virtual router.
<i>vr_name</i>	Specifies the virtual router ID.
<i>local</i>	Specifies the remote Syslog server facility: "local0" "local1" "local2" "local3" "local4" "local5" "local6" "local7".
reference-identifier	Remote Syslog server certificate reference identifier.
<i>reference_identifier</i>	Identifier value (for example, the host name). If none is specified, the existing reference identifier configuration is removed.

Default

If a TLS port is not specified, the default is 6514.

Example

The following example specifies the reference identifier as "hostname" for all specified targets on VR "vr1":

```
# configure syslog all vr vr1 reference-identifier hostname
```

History

This command was first available in ExtremeXOS 22.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure system ports notation

```
configure system ports notation [slot:port | slot/port]
```

Description

Configures a standalone switch to be addressed with a slot number.

Syntax Description

system	Configures system settings.
ports	Configures system ports settings.
notation	Configures system port notation settings.
slot:port	Designates slot:port notation. For example, 1:47. (Default on stacks and Extended Edge Switching). Also designates a 5720 slot:port:channel notation. For example, 5:49:3 for Slot 5 Port 49 Channel 3.
slot	Designates standard slot notation. For example, 1/47.
port	Designates standard port notation. For example, 1/47. Also designates a channelized 5720 port. For example, Port 49 Channel 3.

Default

By default, on standalone switches, port notation is used.

By default, on stacks/Extended Edge Switching, slot:port notation is used.

Usage Guidelines

You can configure a standalone system as a slotted system with this command, which allows for commands which had 'slot' arguments to be visible and take in a valid slot number of '1', along with any port arguments specified in 'slot':port' notation. In turn, any command output would specify 'slot' information and ports displayed in 'slot':port' or slot/port notation.

ExtremeSwitching 5720 series switches use a specified channelized port number.



Note

5720 VIM-6YE ports are not channelized and map to ports 51-53 and 54-56.



Note

Switches running Release 31.6 or earlier that are connected to a channelized 5720 port will not display the correct port number via the Extreme Discovery Protocol. The port numbers will display correctly via the Link Layer Discovery Protocol.

This command requires a configuration save and reboot to take effect.

To view the port notation status, use the `show management` command.

Example

The following example changes a standalone switch to have slot:port notation:

```
# configure system ports notation slot:port
This command will take effect after the next reboot.
```

History

This command was first available in ExtremeXOS 30.2.

The **slot/port** keyword was added in ExtremeXOS 31.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure sys-recovery-level switch

```
configure sys-recovery-level switch [none | reset | shutdown]
```

Description

Configures a recovery option for instances where a hardware exception occurs on ExtremeSwitching series switches.

Syntax Description

none	Configures the switch to maintain its current state regardless of the detected fault. The switch does not reboot or shutdown. ExtremeXOS logs fault and error messages to the syslog.
reset	Configures the switch to reboot upon detecting a hardware fault. ExtremeXOS logs fault, error, system reset, and system reboot messages to the syslog.
shutdown	Configures the switch to shut down upon detecting a hardware fault. All ports are taken offline in response to the reported errors; however, the management port remains operational for debugging purposes only. If the switch shuts down, it remains in this state across additional reboots or power cycles until you explicitly clear the shutdown state.

Default

The default setting is reset.

Usage Guidelines

Use this command for system auto-recovery upon detection of hardware problems. You can configure ExtremeSwitching series switches to take no action, automatically reboot, or shutdown if the switch detects a hardware fault. This enhanced level of recovery detects faults in the CPU.

You must specify one of the following parameters for the switch to respond to hardware failures:

- **none**—Configures the switch to maintain its current state regardless of the detected fault. The switch does not reboot or shutdown.
- **reset**—Configures the switch to reboot upon detecting a hardware fault.
- **shutdown**—Configures the switch to shutdown upon fault detection. All ports are taken offline in response to the reported errors; however, the management port remains operational for debugging purposes only.

Messages Displayed

If you configure the hardware recovery setting to either none (ignore) or shutdown, the switch prompts you to confirm this action by displaying a message similar to the following:

```
Are you sure you want to shutdown on errors? (y/n)
```

Enter y to confirm this action and configure the hardware recovery level. Enter n or press [Enter] to cancel this action.

Displaying the Hardware Recovery Setting

To display the hardware recovery setting, use the following command:

```
show switch
```

If you change the hardware recovery setting from the default (reset) to either none (ignore) or shutdown, the Recovery Mode output is expanded to include a description of the hardware recovery mode. If you keep the default behavior or return to reset, the Recovery Mode output lists only the software recovery setting.

The following truncated output from a ExtremeSwitching series switch displays the software recovery and hardware recovery settings (displayed as Recovery Mode):

```
SysName:          TechPubs Lab
SysLocation:
SysContact:       support@extremenetworks.com, +1 888 257 3000
System MAC:       00:04:96:1F:A5:71
Recovery Mode:    All, Ignore
System Watchdog:  Enabled
```

If you configure the hardware recovery setting to none, the output displays “Ignore” to indicate that no corrective actions will occur on the switch. “Ignore” appears only if you configure the hardware recovery setting to none.

If you configure the hardware recovery setting to shutdown, the output displays “Shutdown” to indicate that the switch will shutdown if fault detection occurs. “Shutdown” appears only if you configure the hardware recovery setting to shutdown.

If you configure the hardware recovery setting to reset, the output displays only the software recovery mode.

Example

The following command configures the switch to not take an action if a hardware fault occurs:

```
# configure sys-recovery-level switch none
```

History

This command was first available in ExtremeXOS 11.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure sys-recovery-level

```
configure sys-recovery-level [all | none]
```

Description

Configures a recovery option for instances where a software exception occurs in ExtremeXOS.

Syntax Description

all	Configures ExtremeXOS to log an error into the syslog and reboot the system after any software task exception occurs.
none	<p>Configures the recovery level to none. No action is taken when a software task exception occurs; there is no system reboot, which can cause unexpected switch behavior.</p> <p>Note: Use this parameter only under the guidance of Extreme Networks Technical Support personnel.</p>

Default

The default setting is all.

Usage Guidelines

If the software fails, the switch automatically reboots or leaves the system in its current state. You must specify one of the following parameters for the system to respond to software failures:

- all—The system will send error messages to the Syslog and reboot if any software task exception occurs.

- none—No action is taken when a software task exception occurs. The system does not reboot, which can cause unexpected switch behavior.

**Note**

Use the none parameter only under the guidance of Extreme Networks Technical Support personnel.

The default setting and behavior is all. Extreme Networks strongly recommends using the default setting.

Displaying the System Recovery Setting

To display the software recovery setting on the switch, use the following command:

```
# show switch
```

This command displays general switch information, including the software recovery level. The following truncated output from an ExtremeSwitching switch displays the software recovery setting (displayed as Recovery Mode):

```
SysName:          TechPubs Lab
SysLocation:
SysContact:       support@extremenetworks.com, +1 888 257 3000
System MAC:       00:04:96:20:B4:13
SysHealth check:  Enabled (Normal)
Recovery Mode:    All
System Watchdog:  Enabled
```

**Note**

All platforms display the software recovery setting as Recovery Mode.

Example

The following command configures a switch to not take an action when any software task exception occurs:

```
# configure sys-recovery-level none
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure tacacs priv-lvl

```
configure tacacs priv-lvl [required | optional]
```

Description

Sets the requirement that the privilege level attribute (priv-lvl) must be specified for TACACS priv-lvl authentication to occur.

Syntax Description

priv-lvl	Specifies setting the requirement that the privilege level attribute for authentication to occur.
required	Fails login attempt if priv-lvl attribute is not provided.
optional	Allows login to occur with read-only privilege if priv-lvl is not provided. (default).

Default

By default, the priv-lvl is not required.

Usage Guidelines

Using this command to set the privilege level attribute as **required** does not change any behavior associated with values received in the priv-lvl attribute, only the presence/absence of the attribute.

Example

The following example makes the priv-lvl attribute required for TACACS authentication:

```
# configure tacacs priv-lvl required
```

History

This command was first available in ExtremeXOS 22.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure tacacs server client-ip

```
configure tacacs [primary | secondary] server [ipaddress | hostname]  
                {tcp_port} client-ip ipaddress {vr vr_name}
```

Description

Configures the server information for a TACACS+ authentication server.

Syntax Description

primary	Configures the primary TACACS+ server.
secondary	Configures the secondary TACACS+ server.
<i>ipaddress</i>	The IP address of the TACACS+ server being configured.
<i>hostname</i>	The host name of the TACACS+ server being configured.
<i>tcp_port</i>	The TCP port to use to contact the TACACS+ server.
<i>ipaddress</i>	The IP address used by the switch to identify itself when communicating with the TACACS+ server.
<i>vr_name</i>	Specifies the virtual router on which the client IP is located. Note: User-created VRs are supported only on the platforms listed for this feature in the Switch Engine 32.3 Feature License Requirements document.

Default

TACACS+ uses TCP port 49. The default virtual router is VR-Mgmt, the management virtual router.

Usage Guidelines

Use this command to configure the server information for a TACACS+ server.

To remove a server, use the following command:

```
unconfigure tacacs server [primary | secondary]
```

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 using a virtual router interface of VR-Default:

```
configure tacacs primary server tacacs1 client-ip 10.10.20.35 vr vr-Default
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure tacacs shared-secret

```
configure tacacs [primary | secondary] shared-secret {encrypted  
    encrypted_secret | secret }
```

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 string is already encrypted.
<i>string</i>	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.

The encrypted keyword is primarily for the output of the show configuration command, so the shared secret is not revealed in the command output. Do not use it to set the shared secret.

Example

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

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

History

This command was first available in ExtremeXOS 10.1.

The encrypted keyword was added in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure tacacs timeout

```
configure tacacs timeout seconds
```

Description

Configures the timeout interval for TACAS+ authentication requests.

Syntax Description

<i>seconds</i>	Specifies the number of seconds for authentication requests. Range is 3 to 120 seconds.
----------------	---

Default

The default is 3 seconds.

Usage Guidelines

Use this command to configure the timeout interval for TACACS+ authentication requests.

To detect and recover from a TACACS+ server failure when the timeout has expired, the switch makes one authentication attempt before trying the next designated TACACS+ server or reverting to the local database for authentication. In the event that the switch still has IP connectivity to the TACACS+ server, but a TCP session cannot be established, (such as a failed TACACS+ daemon on the server), failover happens immediately regardless of the configured timeout value.

For example, if the timeout value is set for 3 seconds (the default value), it will take 3 seconds to fail over from the primary TACACS+ server to the secondary TACACS+ server. If both the primary and the secondary servers fail or are unavailable, it takes approximately 6 seconds to revert to the local database for authentication.

Example

The following command configures the timeout interval for TACACS+ authentication to 10 seconds:

```
configure tacacs timeout 10
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure tacacs-accounting server

```
configure tacacs-accounting [primary | secondary] server [ipaddress |  
  hostname] [udp_port] client-ip ipaddress [vr vr_name]
```

Description

Configures the TACACS+ accounting server.

Syntax Description

primary	Configures the primary TACACS+ accounting server.
secondary	Configures the secondary TACACS+ accounting server.
<i>ipaddress</i>	The IP address of the TACACS+ accounting server being configured.
<i>hostname</i>	The host name of the TACACS+ accounting server being configured.
tcp_port	The TCP port to use to contact the TACACS+ server.
<i>ipaddress</i>	The IP address used by the switch to identify itself when communicating with the TACACS+ accounting server.
<i>vr_name</i>	Specifies the virtual router on which the client IP is located. Note: User-created VRs are supported only on the platforms listed for this feature in the Switch Engine 32.3 Feature License Requirements document.

Default

Unconfigured. The default virtual router is *VR-Mgmt*, the management virtual router.

Usage Guidelines

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

To remove a server, use the following command:

```
unconfigure tacacs server [primary | secondary]
```

Example

The following command configures server tacacs1 as the primary TACACS+ accounting server for client switch 10.10.20.35 using a virtual router interface of *VR-Default*:

```
configure tacacs-accounting primary server tacacs1 client-ip 10.10.20.35 vr vr-Default
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure tacacs-accounting shared-secret

```
configure tacacs-accounting [primary | secondary] shared-secret
    {encrypted encrypted_secret | secret }
```

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.
<i>string</i>	The string to be used for authentication.

Default

N/A.

Usage Guidelines

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

The encrypted keyword is primarily for the output of the show configuration command, so the shared secret is not revealed in the command output. Do not use it to set the shared secret.

Example

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

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

History

This command was first available in ExtremeXOS 10.1.

The encrypted keyword was added in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure tacacs-accounting timeout

```
configure tacacs-accounting timeout seconds
```

Description

Configures the timeout interval for TACACS+ accounting authentication requests.

Syntax Description

<i>seconds</i>	Specifies the number of seconds for accounting requests. Range is 3 to 120 seconds.
----------------	---

Default

The default is 3 seconds.

Usage Guidelines

This command configures the timeout interval for TACACS+ accounting authentication requests.

To detect and recover from a TACACS+ accounting server failure when the timeout has expired, the switch makes one authentication attempt before trying the next designated TACACS+ accounting server or reverting to the local database for authentication. In the event that the switch still has IP connectivity to the TACACS+ accounting server, but a TCP session cannot be established, (such as a failed TACACS+ daemon on the accounting server), failover happens immediately regardless of the configured timeout value.

For example, if the timeout value is set for 3 seconds (the default value), it takes 3 seconds to fail over from the primary TACACS+ accounting server to the secondary TACACS+ accounting server. If both the primary and the secondary servers fail or are unavailable, it takes approximately 6 seconds to revert to the local database for authentication.

Example

The following command configures the timeout interval for TACACS+ accounting authentication to 10 seconds:

```
configure tacacs-accounting timeout 10
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure tech-support add collector

```
configure tech-support add collector [hostname | ip_address] tcp-port  
port {vr vr_name} {from source_ip_address} {ssl [on | off]}
```

Description

This command adds collectors that the switch attempts to connect to for the purpose of forwarding status reports. The collector is identified by its hostname or IP address.

This command also configures the initial value of the TCP port that the collector is listening to, the VR name and source IP address that the switch uses to attempt to connect to the collector, and the SSL mode whether the switch needs to turn SSL on or off when it connects to the collector.

Syntax Description

<i>hostname</i>	Host name of the collector.
<i>ip_address</i>	IPv4 address of the collector.
tcp-port	TCP port number that the collector is listening.
<i>port</i>	Port number. The range is 1-65535.
vr <i>vr_name</i>	Specifies the Virtual router and virtual router name. The default name is <i>VR-Mgmt</i> .
from <i>source_ip_address</i>	Specifies the source and the source IPv4 address. The default source is the IP address on <i>VLAN Mgmt</i> .
ssl	Specifies the Secure Sockets Layer.
on	Specifies that SSL is on.
off	Specifies that SSL is off.

Default

Disabled.

Usage Guidelines

This command adds collectors that the switch attempts to connect to for the purpose of forwarding status reports. The collector is identified by its hostname or IP address. Each added collector needs to have a unique hostname or IP address. If the specified hostname or IP address has already existed, an error message *'ERROR: The collector 1.1.1.1 already exists'* is displayed. Other commands use hostname or IP address to specify the collector that the command reconfigures, deletes, runs reports for, or shows configuration and status.

This command also configures the initial value of the TCP port that the collector is listening to, the VR name and source IP address that the switch uses to connect to the collector, and the SSL mode that determines if the switch needs to turn SSL on/off when connecting to the collector. The purpose of having a default collector configured is to minimize the configuration required for a customer to enable techSupport.

Example

The following command adds a collector at address "1.1.1.1" listening to TCP port "1":

```
configure tech-support add collector 1.1.1.1 tcp-port 1
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure tech-support collector

```
configure tech-support collector [hostname | ip_address] tcp-port port
    {vr vr_name} {from source_ip_address} {ssl [on | off]}
```

Description

This command reconfigures the TCP port, the VR, the Source IP Address, and SSL mode of an existing collector.

Syntax Description

<i>hostname</i>	Host name of the collector.
<i>ip_address</i>	IPv4 address of the collector.
tcp-port	TCP port number that the collector is listening.
<i>port</i>	Port number. The range is 1-65535.
vr <i>vr_name</i>	Specifies the Virtual router and virtual router name. The default name is <i>VR-Mgmt</i> .
from <i>source_ip_address</i>	Specifies the source and the source IPv4 address. The default source is the IP address on <i>VLAN</i> Mgmt.
ssl	Specifies the Secure Sockets Layer.
on	Specifies that SSL is on.
off	Specifies that SSL is off.

Default

Disabled.

Usage Guidelines

This command reconfigures the TCP port, the VR, the Source IP Address, and SSL mode of an existing collector. The collector to be reconfigured is specified by its hostname or IP address. If the specified collector does not exist, an error message `ERROR: The collector 1.1.1.1 does not exists` is displayed.

Example

The following command reconfigures the tech support collector:

```
configure tech-support collector
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure tech-support collector data-set

```
configure tech-support collector [ all hostname | ip_address] data-set
    [ summary | detail ]
```

Description

This command configures the amount and type of data that is included in the status report for a collector.

Syntax Description

all	Configures report data set for all existing collectors.
<i>hostname</i>	Specifies the host name of the collector.
<i>ip_address</i>	Specifies the IPv4 address of the collector.
data-set	Specifies the report data set. The default is detail .
detail	Specifies the output of <code>show tech-support area</code> for area general, config, log, <u>VLAN</u> , and EPM.
summary	Specifies the output of <code>show tech-support all</code> command.

Default

The default is **detail**.

Usage Guidelines

This command configures the amount and type of data that is included in the status report for a collector. When you specify **all**, it configures a report data set for all existing collectors; otherwise report data is set for a particular collector specified by the *hostname* or *IP address*. When the data set is set to **summary**, the status report sent by the switch includes installed ExtremeXOS and Bootrom image versions, the active partition, serial number, equipment type, installed hardware options, stored SRAM contents, basic switch configuration, and log messages. The output of the summary option is

collected from the `show tech-support area` command for the area general, configuration, log, VLAN, and EPM. Changing the report data set to **detail** will send the full output of the `show tech` command. When a collector is added, the data set is set to **detail**.

Example

The following command example configures a specific collector to display a detailed output set:

```
configure tech-support collector 65.222.234.14 data-set detail
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure tech-support collector frequency error-detected

```
configure tech-support collector [ all | hostname | ip_address ]
    frequency [ bootup [ on | off ] | error-detected [ on | off ] | daily [ on
    { time hour } | off ] ]
```

Description

This command configures how often the switch sends status reports for a collector.

Syntax Description

all	Configures report mode for all report collectors.
<i>hostname</i>	Configures report mode based on the host name of the collector.
<i>ip_address</i>	Configures report mode based on the IPv4 address of the collector.
bootup	Send status report when the switch boots up. The default value is on.
on	Specifies that the status reporting is on at bootup.
off	Specifies that the status reporting is off at bootup.
error-detected	Specifies that a status report is sent when a critical severity event is logged. The default value is off.
on	Specifies that error-detected reporting is on.
off	Specifies that error-detected reporting is off. This is the default value.
daily	Specifies that status reports are sent once a day. The default value is off.
on <i>time</i> <i>hour</i>	Specifies the time to send the report. Specifies the hour 0-23. The default value is 0 (12:00AM).
off	Specifies that the daily status reports are off. This is the default value.

Default

Disabled.

Usage Guidelines

This command configures the frequency that the switch sends status reports for a collector. By specifying **all**, it configures report frequency for all existing collectors; otherwise it configures report frequency for a particular collector specified by the hostname or IP address. If the **bootup** option is set to **on**, the switch sends a status report when the switch boots up. If the **error-detected** option is set to **on**, the switch sends a status report when a critical severity event is logged. If the **daily** option is set to **on**, the switch sends a status report once a day regardless of the switch's operational status during the last 24 hour period.

Optionally, you can specify the hour that the report is sent. The default hour is 0, and the valid range is 0 to 23, where 0 is 12:00 AM local time and 23 is 11:00 PM local time. You can enable or disable each option (**bootup**, **error-detected** or **daily**) independently. When all three options of a collector are turned off, the switch does not send any status report to that collector even if the report mode of the collector is set to automatic. When a collector is added, the **bootup** option is set to on, and the **error-detected** and **daily** option is set to off.

Example

The following command example configures the report mode on all existing collectors:

```
configure tech-support collector all report
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure tech-support collector report

```
configure tech-support collector [hostname | ip_address] report
[ automatic | manual ]
```

Description

This command configures the report mode for a collector.

Syntax Description

all	Configures report mode for all report collectors.
<i>hostname</i>	Configures report mode based on the host name of the collector.

<i>ip_address</i>	Configures report mode based on the IPv4 address of the collector.
automatic	Automatically reports switch status to the configured collector (Default).
manual	Manually reports switch status to the configured collector through the <code>run tech-support report</code> .

Default

Disabled.

If enabled, the automatic collector is the default report setting.

Usage Guidelines

This command configures the report mode for a collector. When you specify **all**, it configures report mode for all existing collectors, otherwise it configures report mode for a particular collector specified by the hostname, or IP address. When the report mode is set to **automatic**, the switch automatically attempts to connect to the cloud-hosted collector, and reports the switch status information based on the frequency and data set setting of the collector. Changing the configuration to **manual** restricts reporting to user initiated mode using the `run tech-support` command for that collector. When a collector is added, the report mode is set to automatic by default.

Example

The following command example configures the report mode on all existing collectors:

```
configure tech-support collector all report
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure tech-support delete collector

```
configure tech-support delete collector [ all | hostname | ip_address ]
```

Description

This command deletes existing collectors.

Syntax Description

all	Specifies that you delete all report collectors.
<i>hostname</i>	Specifies the host name of the collector you want to delete.
<i>ip_address</i>	Specifies the IPv4 address of the collector you want to delete.

Default

Disabled.

Usage Guidelines

This command deletes existing collectors. If you specify **all**, it deletes all existing collectors; otherwise it deletes the collector specified by the hostname or IP address. If the specified collector does not exist, an error message `ERROR: The collector 1.1.1.1 does not exist` is displayed.

Example

The following example deletes all collectors :

```
configure tech-support delete collector all
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure telnet access-profile

```
configure telnet access-profile [ access_profile | [[add rule ] [first |  
[before | after] previous_rule]] ] | delete rule | none ]
```

Description

Configures Telnet to use an ACL policy or ACL rule for access control.

Syntax Description

<i>access_profile</i>	Specifies an ACL policy.
add	Specifies that an ACL rule is to be added to the Telnet application.
<i>rule</i>	Specifies an ACL rule.
first	Specifies that the new rule is to be added before all other rules.

before	Specifies that the new rule is to be added before a previous rule.
after	Specifies that the new rule is to be added after a previous rule.
<i>previous_rule</i>	Specifies an existing rule in the application.
delete	Specifies that one particular rule is to be deleted.
none	Specifies that all the rules or a policy file is to be deleted.

Default

Telnet is enabled with no ACL policies and uses TCP port 23.

Usage Guidelines

You must be logged in as administrator to configure Telnet parameters.

You can restrict Telnet access in the following ways:

- Implement an ACL policy file that permits or denies a specific list of IP addresses and subnet masks for the Telnet port. You must create the ACL policy file before you can use this command. If the ACL policy file does not exist on the switch, the switch returns an error message indicating that the file does not exist.

In the ACL policy file for Telnet, the “source-address” field is the only supported match condition. Any other match conditions are ignored.

Use the none option to remove a previously configured ACL.

- Add an ACL rule to the Telnet application through this command. Once an ACL is associated with Telnet, all the packets that reach a Telnet module are evaluated with this ACL and appropriate action (permit or deny) is taken, as is done using policy files.

The permit or deny counters are also updated accordingly regardless of whether the ACL is configured to add counters. To display counter statistics, use the `show access-list counters process telnet` command.

Only the following match conditions and actions are copied to the client memory. Others that may be in the rule are not copied.

Match conditions:

- Source-address—IPv4 and IPv6
- Actions—Permit or Deny

When adding a new rule, use the first, before, and after *previous_rule* parameters to position it within the existing rules.

If the Telnet traffic does not match any of the rules, the default behavior is deny. To permit Telnet traffic that does not match any of the rules, add a permit all rule at the end of the rule list.

Creating an ACL Policy File

To create an ACL policy file, use the `edit policy` command. For more information about creating and implementing ACL policy files, see the *Policy Manager* and *ACLs* chapters in the [Switch Engine 32.3 User Guide](#).

If you attempt to implement a policy that does not exist on the switch, an error message similar to the following appears:

```
Error: Policy /config/MyAccessProfile.pol does not exist on file system
```

If this occurs, make sure the policy you want to implement exists on the switch. To confirm the policies on the switch, use the `configure snmp add community` command. If the policy does not exist, create the ACL policy file.

Viewing Telnet Information

To display the status of Telnet, including the current TCP port, the virtual router used to establish a Telnet session, and whether ACLs are controlling Telnet access, use the following command: `show management`.

Example

The following example applies the ACL policy `MyAccessProfile_2` to Telnet:

```
configure telnet access-profile MyAccessProfile_2
```

The following example applies the ACL rule `DenyAccess` to the Telnet application in the first position in the list:

```
configure telnet access-profile add DenyAccess first
```

The following example removes the association of a single ACL rule from the Telnet application:

```
configure telnet access-profile delete DenyAccess
```

The following example removes the association of an ACL policy or all ACL rules from the Telnet application:

```
configure telnet access-profile none
```

History

This command was first available in ExtremeXOS 11.2.

Support for ACL rules for Telnet was added in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure telnet port

```
configure telnet port [portno | default]
```

Description

Configures the TCP port used by Telnet for communication.

Syntax Description

<i>portno</i>	Specifies a TCP port number. The default is 23. The range is 1 through 65535. The following TCP port numbers are reserved and cannot be used for Telnet connections: 22, 80, and 1023.
default	Specifies the default Telnet TCP port number. The default is 23.

Default

The switch listens for Telnet connections on Port 23.

Usage Guidelines

You must be logged in as administrator to configure the Telnet port.

The portno range is 1 through 65535. The following TCP port numbers are reserved and cannot be used for Telnet connections: 22, 80, and 1023. If you attempt to configure a reserved port, the switch displays an error message similar to the following:

```
configure telnet port 22
Error: port number is a reserved port
```

If this occurs, select a port number that is not a reserved port.

The switch accepts IPv6 connections.

Example

The following command changes the port used for Telnet to port 85:

```
configure telnet port 85
```

The following command returns the port used for Telnet to the default port of 23:

```
configure telnet port default
```

History

This command was first available in ExtremeXOS 10.1.

Support for IPv6 connections was added in ExtremeXOS 11.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure telnet vr

```
configure telnet vr [all | default | vr_name]
```

Description

Configures the virtual router used on the switch for listening for Telnet connections.

Syntax Description

all	Specifies to use all virtual routers for Telnet connections.
default	Specifies to use the default virtual router for Telnet connections. The default router is <i>VR-Mgmt</i> .
<i>vr_name</i>	Specifies the name of the virtual router to use for Telnet connections. NOTE: User-created VRs are supported only on the platforms listed for this feature in the Switch Engine 32.3 Feature License Requirements document.

The default is all.

Usage Guidelines

You must be logged in as administrator to configure the virtual router.

The switch accepts IPv6 connections.

If you specify all, the switch listens on all of the available virtual routers for Telnet connections.

The vr_name specifies the name of the virtual router to use for Telnet connections.

If you specify a virtual router name that does not exist, the switch displays an error message similar to the following:

```
configure telnet vr vr-ttt ^ %% Invalid input detected at '^' marker.
```

Example

The following command configures the switch to listen for and receive Telnet requests on all virtual routers:

```
configure telnet vr all
```

History

This command was first available in ExtremeXOS 11.0.

Support for IPv6 was added in ExtremeXOS 11.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure time

```
configure time month day year hour min sec
```

Description

Configures the system date and time.

Syntax Description

<i>month</i>	Specifies the month. The range is 1-12.
<i>day</i>	Specifies the day of the month. The range is 1-31.
<i>year</i>	Specifies the year in the YYYY format. The range is 2003 to 2036.
<i>hour</i>	Specifies the hour of the day. The range is 0 (midnight) to 23 (11 pm).
<i>min</i>	Specifies the minute. The range is 0-59.
<i>sec</i>	Specifies the second. The range is 0-59.

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. You cannot set the year earlier than 2003 or past 2036. You have the choice of inputting the entire time/date string. If you provide one item at a time and press [Tab], the screen prompts you for the next item. Press [cr] to complete the input.

Example

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

```
configure time 02 15 2002 08 42 55
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure timezone

```
configure timezone {name tz_name} 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

<i>tz_name</i>	Specifies an optional name for this timezone specification. May be up to six alphabetic characters in length. The default is an empty string.
<i>GMT_offset</i>	Specifies a Greenwich Mean Time (GMT) offset, in + or - minutes.
autodst	Enables automatic Daylight Saving Time.
<i>dst-timezone-ID</i>	Specifies an optional name for this DST specification. May be up to six characters in length. The default is an empty string.
<i>dst_offset</i>	Specifies an offset from standard time, in minutes. Value is in the range of 1 to 60. Default is 60 minutes.
<i>floatingday</i>	Specifies the day, week, and month of the year to begin or end DST each year. Format is: <i>week day month</i> where: <i>week</i> is specified as [first second third fourth last] or 1-5. <i>day</i> is specified as [sunday monday tuesday wednesday thursday friday saturday] or 1-7 (where 1 is Sunday). <i>month</i> is specified as [january february march april may june july august september october november december] or 1-12. Default for beginning is second sunday march; default for ending is first sunday november.
<i>absoluteday</i>	Specifies a specific day of a specific year on which to begin or end DST. Format is: <i>month day year</i> where: <i>month</i> is specified as 1-12. <i>day</i> is specified as 1-31. <i>year</i> is specified as 2003-2035. The year must be the same for the begin and end dates.
<i>time_of_day</i>	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 second Sunday in March, and ending every first Sunday in November.

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 second Sunday in March at 2:00 AM and ends the first Sunday in November at 2:00 AM, applies to most of North America (beginning in 2007), and can be configured with the following syntax:
`configure timezone GMT_offset autodst.`

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 (for example, 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 second Sunday in March.
- The **ends** specification defaults to every first Sunday in November.
- 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_offset noautodst.`

Greenwich Mean Time offsets

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. [configure timezone](#) on page 1441 describes the GMT 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	Cape Verde Islands
-2:00	-120	AT - Azores	Azores

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

GMT Offset in Hours	GMT Offset in Minutes	Common Time Zone References	Cities
+13:00	+780	PHOT-Phoenix Island Time	Kanton Island
+14:00	+840	LINT-Line Islands Time	Kiritimati

For name creation guidelines and a list of reserved names, see *Object Names* in the [Switch Engine 32.3 User Guide](#).

Example

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

```
# configure 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:

```
3 configure timezone name EST -300 autodst name EDT 60 begins every second sunday march
at 2 ends every first sunday november at 2:00
# configure timezone name EST -300 autodst name EDT 60 begins every 1 1 4 at 2:00 ends
every 5 1 10 at 2:00
# configure timezone name EST -300 autodst name EDT
# configure timezone -300 autodst
```

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

```
# configure 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:

```
# configure 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 ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure trusted-ports trust-for dhcp-server

```
configure trusted-ports [ports|all] trust-for dhcp-server
```


Description

Configures one or more trusted *DHCP* ports.

Syntax Description

<i>ports</i>	Specifies one or more ports to be configured as trusted ports.
all	Specifies all ports to be configured as trusted ports.

Default

N/A.

Usage Guidelines

To configure trusted DHCP ports, you must first enable DHCP snooping on the switch. To enable DHCP snooping, use the following command:

```
enable ip-security dhcp-snooping {vlan} vlan_name ports [all |  
ports] violation-action [drop-packet {[block-mac | block-port]  
[durationduration_in_seconds | permanently] | none}]] {snmp-trap}
```

Trusted ports do not block traffic; rather, the switch forwards any DHCP server packets that appear on trusted ports. Depending on your DHCP snooping configuration, the switch drops packets and can disable the port temporarily, disable the port permanently, blackhole the MAC address temporarily, blackhole the MAC address permanently, and so on.

If you configure one or more trusted ports, the switch assumes that all DHCP server packets on the trusted port are valid.

Displaying DHCP Trusted Server Information

To display the DHCP snooping configuration settings, including DHCP trusted ports if configured, use the following command: `show ip-security dhcp-snooping {vlan} vlan_name`

To display any violations that occur, including those on DHCP trusted ports if configured, use the following command: `show ip-security dhcp-snooping violations {vlan} vlan_name`

Example

The following command configures ports 2:2 and 2:3 as trusted ports:

```
configure trusted-ports 2:2-2:3 trust-for dhcp-server
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure trusted-servers add server

```
configure trusted-servers [dynamic vlan_id | {vlan} vlan_name] add server  
    ip_address trust-for dhcp-server
```

Description

Configures and enables a trusted *DHCP* server on the switch.

Syntax Description

<i>vlan_name</i>	Specifies the <i>VLAN</i> name.
dynamic	Configuration options for dynamically created VLANs.
<i>vlan_id</i>	VLAN ID tag between 1 and 4,094.
<i>ip_address</i>	Specifies the IP address of the trusted DHCP server.

Default

N/A.

Usage Guidelines

If you configured trusted DHCP server, the switch forwards only DHCP packets from the trusted servers. The switch drops DHCP packets from other DHCP snooping-enabled ports.

You can configure a maximum of eight trusted DHCP servers on the switch.

If you configure a port as a trusted port, the switch assumes that all DHCP server packets on that port are valid.

Displaying DHCP Trusted Server Information

To display the DHCP snooping configuration settings, including DHCP trusted servers if configured, use the following command:

```
show ip-security dhcp-snooping {vlan} vlan_name
```

To display any violations that occur, including those on the DHCP trusted servers if configured, use the following command:

```
show ip-security dhcp-snooping violations {vlan} vlan_name
```

Example

The following command configures a trusted DHCP server on the switch:

```
configure trusted-servers vlan purple add server 10.10.10.10 trust-for dhcp-server
```

History

This command was first available in ExtremeXOS 11.6.

Dynamic VLAN and VLAN ID options added in ExtremeXOS 30.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure trusted-servers delete server

```
configure trusted-servers [dynamic vlan_id | vlan vlan_name] delete
server ip_address trust-for dhcp-server
```

Description

Deletes a trusted *DHCP* server from the switch.

Syntax Description

<i>vlan_name</i>	Specifies the <i>VLAN</i> name.
<i>ip_address</i>	Specifies the IP address of the trusted DHCP server.
dynamic	Configuration options for dynamically created VLANs.
<i>vlan_id</i>	VLAN ID tag between 1 and 4,094.

Default

N/A.

Usage Guidelines

Use this command to delete a trusted DHCP server from the switch.

Displaying DHCP Trusted Server Information

To display the DHCP snooping configuration settings, including DHCP trusted servers if configured, use the following command:

```
show ip-security dhcp-snooping {vlan} vlan_name
```

To display any violations that occur, including those on the DHCP trusted servers if configured, use the following command:

```
show ip-security dhcp-snooping violations {vlan} vlan_name
```

Example

The following command deletes a trusted DHCP server from the switch:

```
configure trusted-servers vlan purple delete server 10.10.10.10 trust-for dhcp-server
```

History

This command was first available in ExtremeXOS 11.6.

Dynamic VLAN and VLAN ID options added in ExtremeXOS 30.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure tunnel ipv6address

```
configure tunnel tunnel_name ipv6address [ipv6-link-local | {eui64}  
      ipv6_address_mask ]
```

Description

Configures an IPv6 address/prefix on a tunnel.

Syntax Description

<i>tunnel_name</i>	Specifies an IPv6 tunnel.
ipv6-link-local	Specifies the link-local address for a tunnel.
eui64	Specifies an EUI64 interface identifier for the lower 64 bits of the address.
<i>ipv6_address_mask</i>	Specifies an IPv6 address / IPv6 prefix length.

Default

N/A.

Usage Guidelines

This command will configure an IPv6 address/prefix route on the specified tunnel.

6to4 tunnels must follow the standard address requirement. The address must be of the form `2002:IPv4_source_endpoint::/16`, where *IPv4_source_endpoint* is replaced by the IPv4 source address of the endpoint, in hexadecimal, colon separated form. For example, for a tunnel

endpoint located at IPv4 address 10.20.30.40, the tunnel address would be 2002:a14:1e28::/16. In hex, 10 is a, 20 is 14, 30 is 1e and 40 is 28.

6in4 tunnels have no restrictions on their address format or prefix allocations.



Note

This command does not work for GRE tunnels. The following error message is displayed:
Error: IPv6 addresses can not be configured on GRE type tunnels!

Example

The following example configures the 6in4 tunnel "link39" with the IPv6 link-local address:

```
configure tunnel link39 ipaddress ipv6-link-local
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 interworking feature in the [Switch Engine 32.3 Feature License Requirements](#) document..

configure tunnel ip tcp adjust-mss

```
configure tunnel tunnel_name ip tcp adjust-mss [off | on tcp_mss_value]
```

Description

Adjusts the TCP Maximum Segment Size (MSS) on GRE Tunnel.

Syntax Description

<i>tunnel_name</i>	Specifies an IPv4 tunnel name.
ip	Specifies IP networking.
tcp	Specifies the TCP protocol.
adjust-mss	Specifies to adjust the TCP payload MSS.
off	Specifies to turn off TCP MSS adjustment (Default).
on	Specifies to turn on TCP MSS adjustment.
<i>tcp_mss_value</i>	Specifies the value of TCP MSS in bytes. Range is 536-9130.

Default

Off.

Usage Guidelines

This command is only available for GRE Tunnel.

For the *tcp_mss_value* option, because it is an adjustment, there is no default value. When TCP MSS is off, the value is ignored.

Example

The following example configures tcp adjust-mss for tunnel "mytunnel" with an adjust-mss value of 1300:

```
configure tunnel mytunnel ip tcp adjust-mss on 1300
```

History

This command was first available in ExtremeXOS 31.6.

Platform Availability

This command is available on the platforms listed for the IPv6 interworking feature in the [Switch Engine 32.3 Feature License Requirements](#) document..

NEW! configure tunnel ipv6 tcp adjust-mss

```
configure tunnel tunnel_name ipv6 tcp adjust-mss [off | on  
    tcp_mss_value]
```

Description

Adjusts the TCP Maximum Segment Size (MSS) on GRE Tunnel.

Syntax Description

<i>tunnel_name</i>	Specifies an IPv6 tunnel name.
ipv6	Specifies IPv6 networking.
tcp	Specifies the TCP protocol.
adjust-mss	Specifies to adjust the TCP payload MSS.
off	Specifies to turn off TCP MSS adjustment (Default).
on	Specifies to turn on TCP MSS adjustment.
<i>tcp_mss_value</i>	Specifies the value of TCP MSS in bytes. Range is 536-9110.

Default

Off.

Usage Guidelines

This command is only available for GRE Tunnel.

For the *tcp_mss_value* option, because it is an adjustment, there is no default value. When TCP MSS is off, the value is ignored.

Example

The following example configures tcp adjust-mss for tunnel "mytunnel" with an adjust-mss value of 1300:

```
configure tunnel mytunnel ipv6 tcp adjust-mss on 1300
```

History

This command was first available in ExtremeXOS 32.3.

Platform Availability

This command is available on the platforms listed for the IPv6 interworking feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure twamp endpoint

```
configure twamp [add | delete] endpoint {vr name} ipaddress ip port  
                        udp_port
```

Description

This command allows you to add and delete the TWAMP endpoints.

Syntax Description

<i>ip</i>	The endpoint IP address, either IPv4 or IPv6.
<i>udp_port</i>	The UDP port the endpoint will listen on; range is 1025 - 65535
<i>name</i>	An optional VR may be used; default is <u><i>VR-Default</i></u> .

Default

N/A.

Usage Guidelines

Use this command to add and delete the TWAMP endpoints. The user specifies the IP address and UDP port number for the endpoint. Removing the endpoint terminates all test sessions associated with the endpoint.

History

This command was first available in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure twamp key-id

```
configure twamp [ add | delete ] key-id key_name shared_secret
```

Description

This command configures the shared secret used for authentication and encryption.

Syntax Description

<i>key_name</i>	The 80 octet KeyID field in the Set-Up-Response control message from RFC 4656.
<i>shared_secret</i>	The shared secret passphrase, which is used to derive the shared secret key, as defined in RFC 5357.

Default

N/A.

Usage Guidelines

None.

History

This command was first available in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure twamp reflector

```
configure twamp reflector [{sessions count} {timeout ref_wait}]
```

Description

This command allows you to modify the number of test sessions to support and timeout value for those test sessions.

Syntax Description

<i>count</i>	Range 0 - 2000 entries; default 2000.
<i>ref_wait</i>	Range 30 - 3600 seconds; default 900 seconds.

Default

count = 2000

ref_wait = 900

Usage Guidelines

The timeout value is the REFWAIT value specified in RFC 5357.

History

This command was first available in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure twamp server

```
configure twamp server [{sessions count} {timeout serv_wait}]
```

Description

This command allows you to modify the number of concurrent TWAMP control sessions to support and the timeout value for those control sessions.

Syntax Description

<i>count</i>	Range 1 - 64.
<i>serv_wait</i>	Range 30 - 3600 seconds.

Default

count = 64

serv_wait = 900

Usage Guidelines

The application terminates the control session if the timeout value expires without the reception of a TWAMP-Control message. This value is the SERVWAIT value specified in RFC 5357.

History

This command was first available in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure upm event

```
configure upm event upm-event profile profile-name ports port_list
```

Description

Configures a pre-defined event that triggers the named profile.

Syntax Description

<i>upm-event</i>	Specifies a pre-defined event type: device-detect, device-undetected, user-authenticate, user-unauthenticated.
<i>profile-name</i>	Specifies the profile to be configured.
<i>port-list</i>	Attaches the UPM profile to the specified port(s).

Default

N/A.

Usage Guidelines

This command configures a profile to be executed when the specified event occurs on the specified port(s).

You can configure multiple user profiles on the same port(s).

Example

The following example shows how to configure a profile on port 1:1, called "profile 1" that is triggered by the event "device-detect":

```
# configure upm event device-detect profile "p1" ports 1:1
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the Universal Port feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure upm profile maximum execution-time

```
configure upm profile profile-name maximum execution-time seconds
```

Description

Defines a maximum execution period for a profile.

Syntax Description

<i>seconds</i>	Defines the execution period in seconds. The range is 2 to 4294967295 seconds.
----------------	--

Default

30 seconds.

Usage Guidelines

If you make a mistake while configuring a profile and the profile loops, it will loop until the end of the maximum execution period. While testing new profiles, consider configuring a relatively short execution time so that any accidental loops do not create long delays during testing.

Example

The following example sets the execution period to 10 seconds:

```
configure upm profile test maximum execution-time 10
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the Universal Port feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure upm timer after

```
configure upm timer timer-name after time-in-secs {every seconds}
```

Description

Creates and names a UPM timer that is activated after the specified time in seconds.

Syntax Description

<i>timer-name</i>	Specifies the name of the UPM timer to be created.
<i>time-in-secs</i>	Configures the interval after which the UPM timer is activated.
<i>seconds</i>	Configures the UPM timer to be activated after every instance of the specified interval.

Default

N/A.

Usage Guidelines

Use this command to configure a timer that activates after the specified time. This is useful for deployment in CLI scripts, because you do not know what the current time will be when the script executes.

When a switch configuration is saved or restored, the UPM timers are activated only at the predetermined timings that were originally configured with the start time.

The periodic timer configured with the every keyword and the one-time timer configured with only the after keyword have a maximum range of one year in seconds (31,622,400 seconds).

Example

The following example configures the UPM timer “A” to be activated every 10 seconds, after an interval of 20 seconds:

```
configure upm timer "timerA" after 20 every 10
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the Universal Port feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure upm timer at

```
configure upm timer timer-name at month day year hour min secs {every  
seconds}
```

Description

Use this command to configure the time setting on a UPM timer.

Syntax Description

<i>timer-name</i>	Specifies the name of the UPM timer to be created.
<i>month</i>	Configures the month when the UPM timer is activated.
<i>day</i>	Configures the day when the UPM timer is activated.
<i>year</i>	Configures the year when the UPM timer is activated.
<i>hour</i>	Configures the hour when the UPM timer is activated.
<i>min</i>	Configures the minute when the UPM timer is activated.
<i>secs</i>	Configures the second when the UPM timer is activated.
<i>seconds</i>	Configures the UPM timer to be activated at every instance of the specified interval.

Default

N/A.

Usage Guidelines

Use this command to when you know the exact time you want an event to execute. If you use this command without the every keyword, the timer is activated once at the specified time. The every keyword configures a periodic timer that is activated at every instance of the time specified in seconds.

When a switch configuration is saved or restored, the UPM timers are activated only at the predetermined timings that were originally configured with the start time.

Example

The following example shows how to configure a timer, T1, that is activated every 10 seconds beginning at 1400 hours on October 16, 2006:

```
# configure upm timer "t1" at 10 16 2006 14 00 00 every 10
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the Universal Port feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure upm timer profile

```
configure upm timer timer-name profile profileName
```

Description

Associates a profile with a UPM timer.

Syntax Description

<i>timer-name</i>	Specifies the name of the UPM timer to be associated with the named profile.
<i>profileName</i>	Specifies the name of the profile to be associated with the UPM timer.

Default

N/A.

Usage Guidelines

Each timer can be attached to only one profile. Once a timer is configured to a profile, it must be unconfigured from that profile before it can be configured to a different profile.

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the Universal Port feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure virtual-network

```
configure virtual-network vn_name [add | delete] [{vlan vlan_name} |  
{vman vman_name | dynamic {vlan} vlan_id}
```

Description

This command adds/removes a tenant VLAN or VMAN to a virtual network.

Syntax Description

<i>vn_name</i>	Alphanumeric string identifying the Virtual Network to be configured.
add	Add a tenant VLAN to the Virtual Network.
delete	Delete a tenant VLAN from the Virtual Network.
vlan	Specifies VLAN.
<i>vlan_name</i>	Name of the tenant VLAN.
vman	Specifies VMAN.
<i>vman_name</i>	Name of the tenant VMAN.
dynamic	Specifies configuring options for dynamically created VLANs. Adds dynamic VLAN's VID to a VNET. You can save this to the configuration and is it persistent across reboots. After reboot when a dynamic VLAN gets created with matching VID, the VLAN is internally applied to the VNET, so that you do not need to reconfigure this every time after reboot.
vlan	Add or delete a tenant VLAN to the Virtual Network
<i>vlan_id</i>	VLAN ID tag between 2 and 4,094.

Default

N/A.

Usage Guidelines

Only a single VLAN/VMAN can be added to a virtual network.

Example

The following example adds a VLAN to an existing virtual network:

```
# configure virtual-network my_virtual_network add vlan vlan100
```

The following example removes a VLAN from an existing virtual network:

```
# configure virtual-network my_virtual_network delete vlan vlan100
```

The following example adds dynamic VLANs with VID 100 to virtual network "my_virtual_network":

```
# configure virtual-network my_virtual_network add dynamic vlan 100
```

History

This command was first available in ExtremeXOS 21.1.

VMAN option added in ExtremeXOS 22.1.

Configuring dynamic VLANs as tenant VLANs was added in ExtremeXOS 30.3.

Platform Availability

This command is supported on the ExtremeSwitching 5320, 5420, 5520, 5720 series switches, and stacks with 5320, 5420, 5520, 5720 slots only.

configure virtual-network add network ports

```
configure virtual-network add network ports [all | portlist]
```

Description

Add ports that can terminate tunnels carrying VXLAN or NVGRE encapsulated traffic.

Syntax Description

add	Add to existing overlay tunnel termination configuration.
network	Configuration related to underlay network.
ports	Select ports that can terminate tunnels carrying VXLAN or NVGRE encapsulated traffic.
all	Select all ports.
<i>portlist</i>	Lists ports to be added.

Default

N/A.

Example

The following example adds ports 1-10 to terminate tunnels carrying VXLAN or NVGRE encapsulated traffic:

```
configure virtual-network add network ports 1-10
```

History

This command was first available in ExtremeXOS 21.1.

Platform Availability

This command is supported on the ExtremeSwitching 5320, 5420, 5520, 5720 series switches, and stacks with 5320, 5420, 5520, 5720 slots only.

configure virtual-network delete network ports

```
configure virtual-network delete network ports [all | portlist]
```


Description

Deletes ports that can terminate tunnels carrying VXLAN or NVGRE encapsulated traffic.

Syntax Description

delete	Delete from existing overlay tunnel termination configuration.
network	Configuration related to underlay network.
ports	Remove ports that can terminate tunnels carrying VXLAN or NVGRE encapsulated traffic.
all	Select all ports.
<i>portlist</i>	Lists ports to be deleted.

Default

N/A.

Example

The following example deletes ports 1-10 to terminate tunnels carrying VXLAN or NVGRE encapsulated traffic:

```
configure virtual-network delete network ports 1-10
```

History

This command was first available in ExtremeXOS 21.1.

Platform Availability

This command is supported on the ExtremeSwitching 5320, 5420, 5520, 5720 series switches, and stacks with 5320, 5420, 5520, 5720 slots only.

configure virtual-network dynamic

```
configure virtual-network dynamic [on | off]
```

Description

Globally controls enabling or disabling auto-creation of virtual networks.

Syntax Description

virtual-network	Virtual overlay network.
dynamic	Configure creation of dynamic virtual networks.

on	Enable creation of dynamic virtual networks by applications such as BGP Auto-peering.
off	Disable creation of dynamic virtual networks by applications such as BGP Auto-peering (default).

Default

By default, automatic creation of virtual networks is disabled.

Usage Guidelines

Creating or deleting BGP Auto-peering enables or disables automatic virtual network creation.

You can view the setting from this command in the `show virtual-network {vn_name | vxlan vni vni | [vlan vlan_name | vman vman_name]}` command.

Example

The following example enables automatic creation of virtual networks:

```
# configure virtual-network dynamic on
```

History

This command was first available in ExtremeXOS 30.3.

Platform Availability

This command is supported on the ExtremeSwitching 5320, 5420, 5520, 5720 series switches, and stacks with 5320, 5420, 5520, 5720 slots only.

configure virtual-network flooding

```
configure virtual-network vn_name flooding [standard | multicast {group
[<grpipaddress> | none]}]
```

Description

This command modifies the flooding mode of an existing virtual network.

Syntax Description

<i>vn_name</i>	Alphanumeric string identifying the Virtual Network to be configured.
standard	Head-end replication to remote endpoints with standard L2 flooding to tenant ports.
multicast	Multicast flooding to remote endpoints with standard L2 flooding to tenant ports.

group	Configure multicast group for flooding of unknown-destination frames (automatically assigned if unspecified).
<i>grpipaddress</i>	IPv4 multicast group address to be used for flooding.
none	Unconfigure multicast group. Flooding changes to auto-assigned multicast group.

Default

Standard.

Usage Guidelines

For auto-assigning multicast groups, you must configure the following command: **configure virtual-network multicast group**.

This command is not allowed on dynamic virtual networks.

Example

To configure multicast flooding mode (group is auto-assigned):

```
configure virtual-network my_virtual_network flooding multicast
```

To configure multicast flooding mode specifying a multicast group:

```
configure virtual-network my_virtual_network flooding multicast group 232.1.1.1
```

To unconfigure multicast group (flooding changes to auto-assigned group):

```
configure virtual-network my_virtual_network flooding multicast none
```

To configure standard flooding mode:

```
configure virtual-network my_virtual_network flooding standard
```

History

This command was first available in ExtremeXOS 31.2.

Platform Availability

This command is supported on the ExtremeSwitching 5320, 5420, 5520, 5720 standalone, and stacks with 5320, 5420, 5520, 5720 slots only.

configure virtual-network local endpoint

```
configure virtual-network local-endpoint [ ipaddress ipaddress { vr vr_name } | none ]
```

Description

This command configures a local IPv4 address to be used as SIP for encapsulated packets.

Syntax Description

ipaddress	Configure the IP address to be used as source IP address for VXLAN packets encapsulated by this gateway.
<i>ipaddress</i>	An existing interface IPv4 address.
vr	VR/VRF instance the IPv4 address is configured on.
<i>vr_name</i>	An existing VR/VRF name.
none	Remove existing IP address configuration for the local tunnel endpoint for this virtual router.

Default

VR-Default.

Usage Guidelines

The address must have been configured as an interface address prior to issuing this command. Although not mandatory, it is strongly recommended that a loopback VLAN IP address be used as the local IP address for tunnels. “VR-Default” is the default for VR/VRF name. ExtremeXOS checks if the given IP address is configured on the VR/VRF. If not configured, the command fails with an appropriate error message. This release of ExtremeXOS supports tunnel termination on a single VR/VRF. That VR/VRF can be a user created. If you intend to change the IP address or the VR/VRF, you can re-issue the same command with a different IP address to effect the change.

Example

To configure a local tunnel endpoint IP address in a user created VR/VRF:

```
configure virtual-network local-endpoint ipaddress 10.10.10.1 vr VR-User
```

To change a local tunnel endpoint to a different IP address within the same VR/VRF:

```
configure virtual-network local-endpoint ipaddress 20.20.20.1 vr VR-User
```

To unconfigure a local tunnel endpoint IP address:

```
configure virtual-network local-endpoint none
```

To change a local tunnel endpoint to a different IP address in a different VR/VRF:

```
configure virtual-network local-endpoint ipaddress 10.10.10.1 vr VR-Default
```

History

This command was first available in ExtremeXOS 21.1.

Platform Availability

This command is supported on the ExtremeSwitching 5320, 5420, 5520, 5720 series switches, and stacks with 5320, 5420, 5520, 5720 slots only.

configure virtual-network monitor

```
configure virtual-network vn_name monitor [ on | off ]
```

Description

Use this command to enable or disable statistics monitoring (byte/packet counters) on a Virtual Network.

Syntax Description

<i>vn_name</i>	Alphanumeric string identifying the Virtual Network to be configured.
on	Enable statistics.
off	Disable statistics.

Default

N/A.

Usage Guidelines

N/A.

Example

To enable statistics monitoring on an existing Virtual Network:

```
configure virtual-network vnet1 monitor on
```

To disable statistics monitoring on an existing Virtual Network

```
configure virtual-network vnet1 monitor off
```

History

This command was first available in ExtremeXOS 21.1.

Platform Availability

This command is supported on the ExtremeSwitching 5320, 5420, 5520, 5720 standalone, and stacks with 5320, 5420, 5520, 5720 slots only.

configure virtual-network multicast group

```
configure virtual-network multicast group [grpipaddress netmask | none]
```

Description

This command is used to configure multicast group range to be used for auto-assigned groups.

Syntax Description

<i>grpipaddress</i>	IPv4 multicast group range for flooding of unknown destination frames
<i>netmask</i>	IPv4 multicast group address netmask (format 255.x.x.x).
none	Unconfigure multicast group range.

Default

Standard.

Usage Guidelines

Example

```
configure virtual-network multicast group 232.1.1.1 255.255.255.255
```

History

This command was first available in ExtremeXOS 31.2.

Platform Availability

This command is supported on the ExtremeSwitching 5320, 5420, 5520, 5720 standalone, and stacks with 5320, 5420, 5520, 5720 slots only.

configure virtual-network name

```
configure virtual-network vn_name name new_name
```

Description

Renames virtual networks.

Syntax Description

virtual-network	Configures virtual networks.
<i>vn_name</i>	Specifies the virtual network to be renamed.
name	Selects renaming the virtual network.
<i>new_name</i>	Specifies the new name for the virtual network.

Default

N/A.

Usage Guidelines

Dynamically created virtual networks are not saved to the configuration. When a dynamically created virtual network is renamed, the virtual network becomes static and is saved to the configuration.

Example

The following example changes the name of the virtual network from "vn1" to "vn2":

```
# configure virtual-network vn name vn2
```

History

This command was first available in ExtremeXOS 30.3.

Platform Availability

This command is supported on the ExtremeSwitching 5320, 5420, 5520, 5720 series switches, and stacks with 5320, 5420, 5520, 5720 slots only.

configure virtual-network remote-endpoint vxlan ipaddress

```
configure virtual-network vn_name [add | delete] remote-endpoint vxlan  
ipaddress ipaddress {vr vr_name}
```

Description

Use this command to add or remove a remote endpoint to a virtual network.

Syntax Description

<i>vn_name</i>	Alphanumeric string identifying the Virtual Network to be configured.
add	Add configuration to the virtual network.
delete	Delete configuration from the virtual network.
<i>ipaddress</i>	A remote endpoint IP address.

vr	VR/VRF instance the remote endpoint is associated with.
<i>vr_name</i>	An existing VR/VRF name.

Default

VR-Default.

Usage Guidelines

This command is only valid when the virtual network is operating in “flooding standard” mode. The remote endpoint will receive unknown destination frames of all types that enter the virtual network from the local endpoint. For “explicit-remotes” flooding mode, the remote endpoints are added when BUM *FDB* entries are added.

Example

To add a remote endpoint to an existing Virtual Network:

```
configure virtual-network my_virtual_network add remote-endpoint vxlan ipaddress 1.2.3.4
```

To remove a remote endpoint from an existing Virtual Network:

```
configure virtual-network my_virtual_network delete remote-endpoint vxlan ipaddress 1.2.3.4
```

History

This command was first available in ExtremeXOS 21.1.

Platform Availability

This command is supported on the ExtremeSwitching 5320, 5420, 5520, 5720 series switches, and stacks with 5320, 5420, 5520, 5720 slots only.

configure virtual-network remote-endpoint vxlan ipaddress monitor

```
configure virtual-network remote-endpoint vxlan ipaddress ipaddress { vr vr_name } monitor [ on | off ]
```

Description

This command enables or disables statistics monitoring (byte/packet counters) on a Virtual Network remote endpoint.

Syntax Description

<i>ipaddress</i>	An existing interface IPv4 address.
vr	VR/VRF instance the IPv4 address is configured on.

<i>vr_name</i>	An existing VR/VRF name.
on	Enable statistics.
off	Disable statistics.

Default

Off.

Usage Guidelines

The command applied on dynamic remote endpoint is not saved to the configuration. If you want it to be saved, convert the remote endpoint to static using the command `create virtual-network remote-endpoint vxlan ipaddress ipaddress {vr vr_name}`.

Example

To enable statistics monitoring on an existing Virtual Network remote endpoint:

```
configure virtual-network remote-endpoint vxlan ipaddress 10.10.10.146 monitor on
```

To disable statistics monitoring on an existing Virtual Network remote endpoint:

```
configure virtual-network remote-endpoint vxlan ipaddress 10.10.10.146 monitor on
```

History

This command was first available in ExtremeXOS 21.1.

Platform Availability

This command is supported on the ExtremeSwitching 5320, 5420, 5520, 5720 series switches, and stacks with 5320, 5420, 5520, 5720 slots only.

configure virtual-network replication-role

```
configure virtual-network replication-role [rnve | replicator | leaf  
      {selected-replicator-default ipaddress}]
```

Description

Configures the assisted replication role on a switch. In addition, if the switch role is leaf, this command also configures the default selected replicator for all virtual networks created afterwards.

Syntax Description

virtual-network	Designates changing the virtual overlay network configuration.
replication-role	Sets the replication role used for VXLAN replication of unknown unicast, broadcast, and multicast traffic.

rnve	Sets regular network virtualization edge node (RNVE) to use head-end replication (default).
replicator	Sets replicate tunneled broadcast and multicast traffic to leaf nodes.
leaf	Sets tunnel broadcast and multicast traffic to replicator.
selected-replicator-default	Sets the default target to send broadcast and multicast traffic to perform replication.
<i>ipaddress</i>	Sets the IPv4 address of a remote tunnel endpoint configured as a replicator.

Default

If not specified, the replicator role is set to RNVE.

Usage Guidelines

The **selected-replicator-default** value is only applied to currently configured virtual networks if the role is changing to **leaf**. It is always applied to virtual networks you create afterwards. To change the **selected-replicator** for already configured virtual networks, use the command `configure virtual-network vn_name selected-replicator [ipaddress ipaddress | none]`.

Do not configure a switch as the new replicator until the prior replicator is unconfigured because packet duplication might occur.

To remove configuration of the assisted replication feature, set the replication role to **rnve**, and the switch reverts back to head-end replication.

Example

The following example configures a node as a replicator:

```
# configure virtual-network replication-role replicator
```

The following example configures a node as a leaf, assigning a default selected replicator for future created virtual networks only:

```
# configure virtual-network replication-role leaf selected-replicator-default 1.2.3.4
```

History

This command was first available in ExtremeXOS 31.1.

Platform Availability

This command is supported on the ExtremeSwitching 5320, 5420, 5520, 5720 series switches, and stacks with 5320, 5420, 5520, 5720 slots only.

configure virtual-network selected-replicator

```
configure virtual-network vn_name selected-replicator [ipaddress
ipaddress | none]
```

Description

Sets, or removes, the selected replicator for an already configured virtual network.

Syntax Description

virtual-network	Designates changing the virtual overlay network configuration.
<i>vn_name</i>	Sets the alphanumeric string identifying the virtual network to be configured.
selected-replicator	Target to send broadcast and multicast traffic to perform replication when configured as leaf.
ipaddress	Designates configuring the IP address of a remote tunnel-endpoint to be used as replicator.
<i>ipaddress</i>	Sets the IPv4 address of the remote tunnel-endpoint configured as a replicator.
none	Removes an existing default selected replicator for this virtual network.

Default

N/A.

Usage Guidelines

Example

The following example configures the remote tunnel-endpoint at the IP address "10.1.1.100" as the selected replicator for the virtual network "vn-blue":

```
# configure virtual-network vn-blue selected-replicator 10.1.1.100
```

The following example removes the selected replicator from the virtual network "vn-blue":

```
# configure virtual-network vn-blue selected-replicator none
```

History

This command was first available in ExtremeXOS 31.1.

Platform Availability

This command is available on the ExtremeSwitching 5520 platform.

configure virtual-network vxlan vni

```
configure virtual-network vn_name vxlan vni [ vni | none ]
```

Description

Use this command to assign a VXLAN VNI to a virtual network.

Syntax Description

<i>vn_name</i>	Alphanumeric string identifying the Virtual Network to be configured.
<i>vni</i>	Virtual Network Identifier value between 1 and 16777215.
none	Remove existing VXLAN VNI configuration for this virtual network.

Default

N/A.

Usage Guidelines

The range of supported VNIs is 1-16777215. The VNI needs to be unique and not more than a one VNI can configured for a virtual-network in this release of ExtremeXOS.

Example

To configure a VXLAN VNI value of 10000 to an existing Virtual Network:

```
configure virtual-network my_virtual_network vxlan vni 10000
```

History

This command was first available in ExtremeXOS 21.1.

Platform Availability

This command is supported on the ExtremeSwitching 5320, 5420, 5520, 5720 series switches, and stacks with 5320, 5420, 5520, 5720 slots only.

configure vlan add nsi | isid

```
configure [{vlan} vlan_name | vlan vlan_id] add [nsi nsi | isid isid]
```

Description

Maps a static VLAN to a Network Service Identifier (NSI) or Individual Service Identifier (ISID).

Syntax Description

vlan	Specifies VLAN.
<i>vlan_name</i>	Specifies the name of the VLAN to map.
<i>vlan_id</i>	Specifies the ID of the VLAN to map.
add	Specifies mapping a VLAN to an NSI.
nsi	Specifies an NSI.
<i>nsi</i>	Specifies the ID number of the NSI to map to the VLAN.
isid	Specifies an ISID.
<i>isid</i>	Specifies the ID number of the ISID to map to the VLAN.

Default

N/A.

Usage Guidelines

These static VLAN mappings do not age out of the *LLDP* database, but are removed when the VLAN is deleted or when removed by the command [configure vlan delete nsi | isid](#) on page 1480.

You can only map one VLAN to an NSI or ISID.

Example

The following example maps VLAN "vlan1" to NSI "1000":

```
# configure vlan vlan1 add nsi 1000
```

History

This command was first available in ExtremeXOS 22.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure vlan add ports

```
configure [ {vlan} vlan_name | vlan vlan_list] add ports [port_list |
all] {tagged tag | untagged} [{stpd} stpd_name] {dot1d | emistp |
pvst-plus}}
```

Description

Adds one or more ports in a *VLAN*.

Syntax Description

<i>vlan_name</i>	Specifies a VLAN name.
<i>vlan_list</i>	Specifies a VLAN list of IDs.
<i>port_list</i>	Specifies a list of ports or slots and ports.
all	Specifies all ports.
tagged <i>tag</i>	Specifies the ports should be configured as tagged.
untagged	Specifies the ports should be configured as untagged.
<i>stpd_name</i>	Specifies an <i>STP</i> domain name.
dot1d emistp pvst-plus	Specifies the BPDU encapsulation mode for these STP ports.

Default

Untagged.

Usage Guidelines

The VLAN must already exist 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. If you attempt to add an untagged port to a VLAN prior to removing it from the default VLAN, you see the following error message:

```
Error: Protocol conflict when adding untagged port 1:2. Either add this port as tagged or
assign another protocol to this VLAN.
```



Note

This message is not displayed if keyword **all** is used as *port_list*.

The ports that you add to a VLAN and the VLAN itself cannot be explicitly assigned to different virtual routers (VRs). When multiple VRs are defined, consider the following guidelines while adding ports to a VLAN:

- A VLAN can belong (either through explicit or implicit assignment) to only one VR.
- If a VLAN is not explicitly assigned to a VR, then the ports added to the VLAN must be explicitly assigned to a single VR.

- If a VLAN is explicitly assigned to a VR, then the ports added to the VLAN must be explicitly assigned to the same VR or to no VR.
- If a port is added to VLANs that are explicitly assigned to different VRs, the port must be explicitly assigned to no VR.

**Note**

User-created VRs are supported only on the platforms listed for this feature in the [Switch Engine 32.3 Feature License Requirements](#) document. On switches that do not support user-created VRs, all VLANs are created in [VR-Default](#) and cannot be moved.

Refer to the STP section in the [Switch Engine 32.3 User Guide](#) for more information on configuring Spanning Tree Domains.

**Note**

If you use the same name across categories (for example, [STPD](#) and [EAPS](#) names), we recommend that you specify the identifying keyword as well as the actual name. If you do not use the keyword, the system may return an error message.

Beginning with ExtremeXOS 11.4, the system returns the following message if the ports you are adding are already EAPS primary or EAPS secondary ports:

WARNING: Make sure Vlan1 is protected by EAPS. Adding EAPS ring ports to a VLAN could cause a loop in the network. Do you really want to add these ports? (y/n)

Example

The following example assigns tagged ports 1:1, 1:2, 1:3, and 1:6 to a VLAN named "accounting":

```
configure vlan accounting add ports 1:1, 1:2, 1:3, 1:6 tagged
```

History

This command was first available in ExtremeXOS 10.1.

The **tagged** keyword was added in ExtremeXOS 15.4.

The `vlan_list` option was added in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure vlan add ports private-vlan translated

Translation from network [VLAN](#) tag to each subscriber VLAN tag is done by default in a private VLAN.

```
configure [ {vlan} vlan_name | vlan vlan_id] add ports port_list
private-vlan translated
```

Description

Adds the specified ports to the specified network VLAN and enables tag translation for all subscriber VLAN tags to the network VLAN tag.

Syntax Description

<i>vlan_name</i>	Specifies the network VLAN to which the ports are added.
<i>vlan_list</i>	Specifies a VLAN list of IDs.
<i>port_list</i>	Specifies the ports to be added to the network VLAN.

Default

N/A.

Usage Guidelines

This command is allowed only when the specified VLAN is configured as a network VLAN on a PVLAN.

Example

The following example adds port 2:1 to VLAN sharednet and enables VLAN translation on that port:

```
configure sharednet add ports 2:1 private-vlan translated
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms that support the Private VLAN feature. For features and the platforms that support them, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure vlan add ports stpd

```
configure vlan vlan_name add ports [all | port_list] {tagged {tag} |  
  untagged} stpd stpd_name {[dot1d | emistp | pvst-plus]}
```

Description

Adds one or more ports in a VLAN to a specified STPD.

Syntax Description

<i>vlan_name</i>	Specifies a VLAN name.
all	Specifies all of the ports to be included in the STPD.

<i>port_list</i>	Specifies the port or ports to be included in the STPD.
tagged	Specifies the ports should be configured as tagged.
<i>tag</i>	Specifies the port-specific VLAN tag. When there are multiple ports specified in the <i>port_list</i> , the same tag is used for all of them. When unspecified port <i>tag</i> is equal to the VLAN tag.
untagged	Specifies the ports should be configured as untagged.
<i>stpd_name</i>	Specifies an STPD name on the switch.
dot1d	Specifies the <i>STP</i> encapsulation mode of operation to be 802.1d.
emistp	Specifies the STP encapsulation mode of operation to be EMISTP.
pvst-plus	Specifies the STP encapsulation mode of operation to be PVST+.

Default

Ports in the default STPD (s0) are in dot1d mode.

Ports in user-created STPDs are in emistp mode.

Usage Guidelines

To create a VLAN, use the `create vlan` command. To create an STP domain, use the `create stpd` command.

In an EMISTP or PVST+ environment, this command adds a list of ports to a VLAN and a specified STPD at the same time provided the carrier VLAN already exists on the same set of ports. You can also specify the encapsulation mode for those ports.

In an *MSTP* environment, you do not need a carrier VLAN. A CIST controls the connectivity of interconnecting MSTP regions and sends BPDUs across the regions to communicate region status. You must use the dot1d encapsulation mode in an MSTP environment.

You cannot configure STP on the following ports:

- Mirroring target ports.
- Software-controlled redundant ports.

If you see an error similar to the following:

```
Error: Cannot add VLAN default port 3:5 to STP domain
```

You might be attempting to add:

- A carrier VLAN port to a different STP domain than the carrier VLAN belongs.
- A VLAN/port for which the carrier VLAN does not yet belong.



Note

This restriction is only enforced in an active STP domain and when you enable STP to ensure you have a legal STP configuration.

Naming Conventions

If your VLAN has the same name as another component, for example an STPD, we recommend that you specify the identifying keyword as well as the name. If your VLAN has a name unique only to that VLAN, the keywords `vlan` and `stpd` are optional.

STP Encapsulation Modes

You can specify the following STP encapsulation modes:

- `dot1d`—This mode is reserved for backward compatibility with previous STP versions. BPDUs are sent untagged in 802.1D mode. Because of this, any given physical interface can have only one STPD running in 802.1D mode.

This encapsulation mode supports the following STPD modes of operation: 802.1D, 802.1w, and MSTP.

- `emistp`—This mode sends BPDUs with an 802.1Q tag having an STPD ID in the VLAN ID field.

This encapsulation mode supports the following STPD modes of operation: 802.1D and 802.1w.

- `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.

This encapsulation mode supports the following STPD modes of operation: 802.1D and 802.1w.

These encapsulation 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.

MSTP STPDs use only 802.1D BPDU encapsulation mode. The switch prevents you from configuring EMISTP or PVST+ encapsulation mode for MSTP STPDs.

Specify the port `tag` when you need to put multiple vlans into a broadcast domain.

Automatically Inheriting Ports--MSTP Only

In an MSTP environment, whether you manually or automatically bind a port to an *MSTI* in an MSTP region, the switch automatically binds that port to the CIST. The CIST handles BPDU processing for itself and all of the MSTIs; therefore, the CIST must inherit ports from the MSTIs in order to transmit and receive BPDUs.

Example

The following command adds slot 1, port 2 and slot 2, port 3, members of a VLAN named Marketing, to the STPD named STPD1, and specifies that they be in EMISTP mode:

```
configure vlan marketing add ports 1:2, 2:3 tagged stpd stpd1 emistp
```

The following examples illustrate the `tag` variable in ExtremeXOS 15.4.

The following example configures vlan with tag 100 and port tag of 10 and 11 on two different ports:

```
create vlan exchange tag 100
config vlan exchange add ports 3 tagged 10
config vlan exchange add ports 4 tagged 11
```

The following example configures a VLAN with tag 100, and port tag of 10 and 11 on the same ports:

```
create vlan exchange tag 100
config vlan exchange add ports 3 tagged 10
config vlan exchange add ports 3 tagged 11
```

The following example configures VLAN with tag 100, and port tag of 10 on two ports and 11 on a different port:

```
create vlan exchange tag 100
config vlan exchange add ports 2:3,2:4 tagged 10
config vlan exchange add ports 2:5 tagged 11
```

History

This command was first available in ExtremeXOS 10.1.

The **nobroadcast** keyword was removed in ExtremeXOS 11.4.

The *tag* variable was added in ExtremeXOS 15.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure vlan add secondary-ipaddress

```
configure [ {vlan} vlan_name | vlan vlan_id] add secondary-ipaddress
anycast [ip_address {netmask} | ipNetmask]
```

Description

Configures secondary IP addresses on a VLAN to support multinetting.

Syntax Description

<i>vlan_name</i>	Specifies a VLAN name.
<i>vlan_id</i>	Specifies a VLAN id.
anycast	Specifies that the address functions like any other secondary address on a VLAN, but responds to ARP requests with the anycast MAC address.
<i>ip_address</i>	Specifies an IP address.
<i>netmask</i>	Specifies a network mask.
<i>ipNetmask</i>	Specifies an IP address with network mask.

Default

N/A.

Usage Guidelines

Adding a secondary IP address to a VLAN enables multinetting. Secondary addresses are added to support legacy stub IP networks.

After you have added a secondary IP address to a VLAN, you cannot unconfigure the primary IP address of that VLAN until you delete all of the secondary addresses. To delete the secondary address, use the following command:

```
configure [ {vlan} vlan_name | vlan vlan_id] delete secondary-ipaddress
[ip_address | all]
```

Example

The following example configures the VLAN multi to support the 10.1.1.0/24 subnet in addition to its primary subnet:

```
# configure vlan multi add secondary-ipaddress 10.1.1.1/24
```

History

This command was first available in ExtremeXOS 11.0.

The *vlan_id* variable is first available in ExtremeXOS 16.1.

Anycast capability was added in ExtremeXOS 30.7.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure vlan delete nsi | isid

```
configure [{vlan} vlan_name | vlan vlan_id] delete [nsi nsi | isid isid]
```

Description

Unmaps a static VLAN from a Network Service Identifier (NSI) or Individual Service Identifier (ISID).

Syntax Description

vlan	Specifies VLAN.
<i>vlan_name</i>	Specifies the name of the VLAN to remove.
<i>vlan_id</i>	Specifies the ID of the VLAN to remove.
delete	Specifies removing the mapping of a VLAN from an NSI.

nsi	Specifies an NSI.
<i>nsi</i>	Specifies the ID number of the NSI to unmap from the VLAN.
isid	Specifies an ISID.
<i>isid</i>	Specifies the ID number of the ISID to unmap from the VLAN.

Default

N/A.

Usage Guidelines

Only mappings created by the command [configure vlan add nsi | isid](#) on page 1472 can be removed using this command.

Example

The following example removes VLAN "vlan1" from NSI "1000":

```
# configure vlan vlan1 delete nsi 1000
```

History

This command was first available in ExtremeXOS 22.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure vlan delete ports

```
configure [ {vlan} vlan_name | vlan vlan_list] delete ports [all |  
    port_list ]
```

Description

Deletes one or more ports in a [VLAN](#).

Syntax Description

<i>vlan_name</i>	Specifies a VLAN name.
<i>vlan_list</i>	Specifies a VLAN list of IDs.
all	Specifies all ports.
<i>port_list</i>	Specifies a list of ports or slots and ports.
tagged <i>tag</i>	Specifies the port-specific VLAN tag. When there are multiple ports specified using <i>port_list</i> , the same tag is used for all of them.

Default

When unspecified, the port tag is equal to the VLAN tag.

Usage Guidelines

Specify port tag to delete a VLAN port that has a different tag from the VLAN tag.

Example

The following example removes ports 1, 3, and 7 on a switch from a VLAN named accounting:

```
configure accounting delete ports 1,3,7
```

The following example deletes a VLAN port with tag 10:

```
create vlan exchange tag 100
config vlan exchange del ports 3 tag 10
```

The following example deletes a VLAN port tag of 10 on two ports:

```
create vlan exchange tag 100
config vlan exchange d ports 3,4 tag 10
```

History

This command was first available in ExtremeXOS 10.1.

The `vlan_list` option was added in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure vlan delete secondary-ipaddress

```
configure [ {vlan} vlan_name | vlan vlan_id] delete secondary-ipaddress
[ip_address | all]
```

Description

Removes secondary IP addresses on a VLAN that were added to support multinetting.

Syntax Description

<i>vlan_name</i>	Specifies a VLAN name.
<i>vlan_id</i>	Specifies a VLAN ID.
<i>ip_address</i>	Specifies an IP address.
all	Specifies deleting all existing secondary and anycast IP addresses.

Default

N/A.

Usage Guidelines

After you have added a secondary IP address to a VLAN (`configure [{vlan} vlan_name | vlan vlan_id]add secondary-ipaddress anycast [ip_address {netmask} | ipNetmask]`), you cannot unconfigure the primary IP address of that VLAN until you delete all of the secondary addresses. Use the **all** keyword to delete all the secondary and anycast IP addresses from a VLAN.

Example

The following example removes the 10.1.1.0 secondary IP address from the VLAN "multi":

```
# configure vlan multi delete secondary-ipaddress 10.1.1.1
```

History

- This command was first available in ExtremeXOS 11.0.
- The `vlan_id` variable is first available in ExtremeXOS 16.1.
- The capability to delete anycast IP addresses was added in ExtremeXOS 30.7.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure vlan description

```
configure {vlan} vlan_name description [vlan-description | none]
```

Description

Configures a description for the specified VLAN.

Syntax Description

<i>vlan_name</i>	Specifies the VLAN name.
<i>vlan-description</i>	Specifies a VLAN description (up to 64 characters) that appears in show vlan commands and can be read from the ifAlias MIB object for the VLAN.
none	This keyword removes the configured VLAN description.

Default

By default, the VLAN has no description.

Usage Guidelines

The VLAN description must be in quotes if the string contains any space characters. If a VLAN description is configured for a VLAN that already has a description, the new description replaces the old description.

Example

The following example assigns the description "Campus A" to VLAN vlan1:

```
configure vlan vlan1 description "Campus A"
```

History

This command was first available in ExtremeXOS 12.4.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure vlan dhcp-address-range

```
configure vlan vlan_name dhcp-address-range ipaddress1 - ipaddress2
```

Description

Configures a set of DHCP addresses for a VLAN.

Syntax Description

<i>vlan_name</i>	Specifies the VLAN on whose ports DHCP will be enabled.
<i>ipaddress1</i>	Specifies the first IP address in the DHCP address range to be assigned to this VLAN.
<i>ipaddress2</i>	Specifies the last IP address in the DHCP address range to be assigned to this VLAN.

Default

N/A.

Usage Guidelines

The following error conditions are checked: ipaddress2 >= ipaddress1, the range must be in the VLAN's network, the range does not contain the VLAN's IP address, and the VLAN has an IP address assigned.

Example

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

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

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure vlan dhcp-lease-timer

```
configure vlan vlan_name dhcp-lease-timer lease-timer
```

Description

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

Syntax Description

<i>vlan_name</i>	Specifies the <i>VLAN</i> on whose ports netlogin should be disabled.
<i>lease-timer</i>	Specifies the timer value, in seconds.

Default

N/A.

Usage Guidelines

The timer value is specified in seconds. The timer value range is 0 - 4294967295, where 0 indicates the default (not configured) value of 7200 second.

Example

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

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

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure vlan dhcp-options

```
configure {vlan} vlan_name dhcp-options [code option_number [16-bit value1 {value2 {value3 {value4}}}] | 32-bit value1 {value2 {value3 {value4}}}] | flag [on | off] | hex string_value | ipaddress ipaddress1 {ipaddress2 {ipaddress3 {ipaddress4}}}] | string string_value] | default-gateway | dns-server {primary | secondary} | 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

<i>vlan_name</i>	Specifies the <u>VLAN</u> on which to configure DHCP.
code	Specifies the generic DHCP option code.
<i>option_number</i>	Specifies the DHCP Option number.
16-bit	Specifies that one to four 16-bit unsigned integer values associated with selected DHCP option.
32-bit	Specifies that one to four 32-bit unsigned integer values associated with selected DHCP option.
flag	Specifies that 1 byte value associated with selected DHCP option number.
hex	Specifies that hexadecimal string associated with selected DHCP option number.
string	Specifies that a string is associated with selected DHCP option number.
<i>string_value</i>	The string value associated with specified option.
default-gateway	Specifies the router option.
dns-server	Specifies the Domain Name Server (DNS) option.
primary	Specifies the primary DNS option.
secondary	Specifies the secondary DNS option.
wins-server	Specifies the NetBIOS name server (NBNS) option.
<i>ipaddress</i>	The IP address associated with the specified option.

Default

N/A.

Usage Guidelines

This command configures the DHCP options that can be returned to the DHCP client. For the default-gateway option you are only allowed to configure an IP address that is in the VLAN's network range. For the other options, any IP address is allowed.

The options below represent the following BOOTP options specified by RFC2132:

- default-gateway—Router option, number 3.
- dns-server—Domain Name Server option, number 6.
- wins-server—NetBIOS over TCP/IP Name Server option, number 44.

Example

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

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

History

This command was first available in ExtremeXOS 11.0.

The primary and secondary DNS options were added in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure vlan dynamic-vlan uplink-ports

```
configure vlan dynamic-vlan uplink-ports [ add {ports} port_list |  
      delete {ports} [port_list | all] ]
```

Description

Statically provisions uplink ports for all dynamically created VLANs.

Syntax Description

dynamic-vlan	Configuration options for dynamically created VLANs.
uplink-ports	Tagged uplink ports for VLANs created by system applications.
add	Add ports to dynamic VLAN uplink port list.
delete	Remove ports from dynamic VLAN uplink port list.
ports	Ports to be configured as uplink ports.
<i>port_list</i>	List of ports separated by a comma or -.
all	Clear the dynamic VLAN uplink port list.

Default

N/A.

Usage Guidelines

Use this command to statically provision uplink ports for dynamically created VLANs.

Example

```
# conf vlan dynamic-vlan uplink-ports add ports 16-18
# conf vlan dynamic-vlan uplink-ports add 20,22,24
# configure vlan dynamic-vlan uplink-ports delete ports 22
# configure vlan dynamic-vlan uplink-ports delete 16-18
# configure vlan dynamic-vlan uplink-ports delete all
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure vlan ipaddress

```
configure [ {vlan} vlan_name | vlan vlan_id] ipaddress
    [ipaddress {netmask} | {ipNetmask} | ipv6-link-local | {eui64}
    ipv6_address_mask]
```

Description

Assigns an IPv4 address and an optional subnet mask or an IPv6 address to the VLAN. Beginning with ExtremeXOS 11.2, you can specify IPv6 addresses. You can assign either an IPv4 address, and IPv6 address, or both to the VLAN. Beginning with ExtremeXOS 11.3, you can use this command to assign an IP address to a specified VMAN and enable multicasting on that VMAN.



Note

You can also use this command to assign an IP address to a VMAN on all platforms that support the VMAN feature. For information on which software licenses and platforms support the VMAN feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

Syntax Description

<i>vlan_name</i>	Specifies a VLAN name.
<i>vlan_id</i>	Specifies a VLAN ID.
<i>ipaddress</i>	Specifies an IPv4 address.

<i>netmask</i>	Specifies an IPv4 subnet mask in dotted-quad notation (for example, 255.255.255.0). This parameter supports 255.255.255.254 for 31-bit prefixes.
<i>ipNetmask</i>	Specifies an IPv4 prefix mask in CIDR notation. This parameter supports /31 for 31-bit prefixes.
ipv6-link-local	Specifies IPv6 and configures a link-local address generated by combining the standard link-local prefix with the automatically generated interface in the EUI-64 format. Using this option automatically generates an entire IPv6 address; this address is only a link-local, or VLAN-based, IPv6 address; that is, ports on the same segment can communicate using this IP address and do not have to pass through a gateway.
eui64	Specifies IPv6 and automatically generates the interface ID in the EUI-64 format using the interface's MAC address. Once you enter this parameter, you must add the following variables: <i>ipv6_address_mask</i> . Use this option when you want to enter the 64-bit prefix and use a EUI-64 address for the rest of the IPv6 address.
<i>ipv6_address_mask</i>	Specify the IPv6 address in the following format: x:x:x:x:x:x/prefix length, where each x is the hexadecimal value of one of the 8 16-bit pieces of the 128-bit wide address.

Default

N/A.

Usage Guidelines

The VLAN must already exist before you can assign an IP address; use the `create vlan` command to create the VLAN (also the VMAN must already exist).



Note

If you plan to use the VLAN as a control VLAN for an *EAP*s domain, do NOT configure the VLAN with an IP address. For information about adding secondary IP addresses to VLANs, see the *IPv4 Unicast Routing* section in the [Switch Engine 32.3 User Guide](#).

Beginning with ExtremeXOS 11.2, you can specify IPv6 addresses. For information about IPv6 addresses, see the *IPv6 Unicast Routing* section in the [Switch Engine 32.3 User Guide](#).

Beginning with ExtremeXOS 11.3, you can assign an IP address (including IPv6 addresses) to a VMAN. Beginning with version 11.4, you can enable multicasting on that VMAN.

Beginning with ExtremeXOS 15.7.1, you can configure IPv4 addresses with 31-bit prefixes on network VLANs and the Mgmt VLAN.

To enable multicasting on the specified VMAN once you assigned an IP address, take the following steps:

1. Enable IP multicast forwarding.
2. Enable and configure multicasting.

Example

The following examples are equivalent; both assign an IPv4 address of 10.12.123.1 to a VLAN named "accounting":

```
configure vlan accounting ipaddress 10.12.123.1/24
configure vlan accounting ipaddress 10.12.123.1 255.255.255.0
```

The following example assigns a link local IPv6 address to a VLAN named management:

```
configure vlan accounting ipaddress ipv6-link-local
```

History

This command was first available in ExtremeXOS 10.1.

The IPv6 parameters were added in ExtremeXOS 11.2.

Support for 31-bit prefixes on IPv4 addresses was added in ExtremeXOS in 15.7.1.

The `vlan_id` variable is first available in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure vlan l2pt profile

```
configure [[{vlan vlan_name} [vxlan {vr vr_name} rtep rtep_ipv4]] l2pt
profile [none | profile_name]
```

Description

Configures VLAN L2PT profiles on service interfaces.

Syntax Description

vlan	Specifies the <u>VLAN</u> configuration.
<i>vlan_name</i>	Specifies the VLAN name.
vxlan	Specifies Virtual eXtensible LAN.
vr	Specifies Virtual Router.
<i>vr_name</i>	Specifies the Virtual Router Name. If not specified, the VR of the current command context is used.
rtep	Specifies Remote Tunnel End Point.
<i>rtep_ipv4</i>	Specifies the Remote Tunnel End Point IPv4 address.
l2pt	Specifies Layer 2 protocol tunneling.
profile	Specifies the L2PT profile for the RTEP.

none	Specifies that no L2PT profile should be bound to the ports (default).
<i>profile_name</i>	Specifies the L2PT profile to be bound to the ports.

Default

Disabled.

Usage Guidelines

Use this command to configure VLAN L2PT profiles on VXLAN RTEP interfaces.

Example

The following example binds the tenant VLAN to *l2pt-nw* profile with RTEP IP address as 2.2.2.2 of VxLAN service Interface with the action "none":

```
# configure l2pt profile "l2pt-nw" add protocol filter cdp action none
# configure vlan tenant vxlan rtep 2.2.2.2 l2pt profile l2pt-nw
```

The following example unbinds the configured l2pt profile from RTEP IP address 2.2.2.2 associated with tenant VLAN:

```
# configure vlan tenant vxlan rtep 2.2.2.2 l2pt profile none
```

The following example binds the tenant with peer 2.2.2.2 of VxLAN RTEP *l2pt-nw* specifies tunneling actions:

```
# configure l2pt profile "l2pt-nw" add protocol filter cdp action tunnel
# configure tenant vxlan rtep 2.2.2.2 l2pt profile "l2pt-nw"
Error: Cannot tunnel on VxLAN RTEP. Tunnel action may be applied only to ports.
```

History

This command was first available in ExtremeXOS 31.3.

Platform Availability

This command is supported on the ExtremeSwitching 5520 series switch and stacks with 5520 slots only.

configure vlan name

```
configure [ {vlan} vlan_name | vlan vlan_id] name name
```

Description

Renames a previously configured VLAN.

Syntax Description

<i>vlan_name</i>	Specifies the current (old) VLAN name.
<i>vlan_id</i>	Specifies the VLAN ID.
<i>name</i>	Specifies a new name for the VLAN.

Default

N/A.

Usage Guidelines

You cannot change the name of the default VLAN “Default.”

For information on VLAN name requirements and a list of reserved keywords, see *Object Names* in the [Switch Engine 32.3 User Guide](#).



Note

If you use the same name across categories (for example, *STPD* and *EAPS* names), we recommend that you specify the identifying keyword as well as the actual name. If you do not use the keyword, the system may return an error message.

Example

The following example renames VLAN vlan1 to engineering:

```
configure vlan vlan1 name engineering
```

History

This command was first available in ExtremeXOS 10.1.

The *vlan_id* variable is first available in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure vlan netlogin-lease-timer

```
configure vlan vlan_name netlogin-lease-timer seconds
```

Description

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

Syntax Description

<i>vlan_name</i>	Specifies the <u>VLAN</u> to which this timer value applies.
<i>seconds</i>	Specifies the timer value, in seconds.

Default

10 seconds.

Usage Guidelines

The timer value is specified in seconds.

This command applies only to the web-based authentication mode of network login.

Example

The following command sets the timer value to 15 seconds for VLAN corp:

```
configure vlan corp netlogin-lease-timer 15
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure vlan qosprofile

```
configure [ {vlan} vlan_name | vlan vlan_list] {qosprofile} [qosprofile  
| none]
```

Description

Configures a VLAN traffic group, which links all the ingress ports in the specified VLAN to the specified egress QoS profile.

Syntax Description

<i>vlan_name</i>	Specifies a VLAN name.
<i>vlan_list</i>	Specifies a VLAN list of IDs.
<i>qosprofile</i>	Specifies an egress QoS profile. The supported values are: qp1 to qp8 and none.

Default

None.

Usage Guidelines

Extreme switches support eight egress QoS profiles (QP1 to QP8) for each port. The QoS profile QP7 is not available to you on a SummitStack.

Example

The following command configures VLAN accounting to use QoS profile QP3:

```
configure vlan accounting qosprofile qp3
```

History

This command was first available in ExtremeXOS 11.0.

The `vlan_list` option was added in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure vlan protocol

```
configure [ {vlan} vlan_name | vlan vlan_list]protocol {filter}  
           filter_name
```

Description

Configures a VLAN to use a specific protocol filter.

Syntax Description

<i>vlan_name</i>	Specifies a VLAN name.
<i>vlan_list</i>	Specifies a VLAN list of IDs.
<i>protocol_name</i>	Specifies a protocol filter name. This can be the name of a predefined protocol filter, or one you define. The following protocol filters are predefined: IP, IPv6, IPX, NetBIOS, DECNet, IPX_8022, IPX_SNAP, AppleTalk. Using 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 `configure protocol` command to define your own protocol filter.

Protocol Filters

These devices do not forward packets with a protocol-based VLAN set to AppleTalk. To ensure that AppleTalk packets are forwarded on the device, create a protocol-based VLAN set to "any" and define other protocol-based VLANs for other traffic, such as IP traffic. The AppleTalk packets pass on the "any" VLAN, and the other protocols pass traffic on their specific protocol-based VLANs.

Example

The following example configures the protocol filter "my_filter" to vlan v1:

```
configure vlan v1 protocol "my_filter"  
configure vlan v1 protocol filter "my_filter"
```

History

This command was first available in ExtremeXOS 10.1.

The IPv6 parameter was added in ExtremeXOS 11.2.

The **filter** keyword was added in ExtremeXOS 15.5.

The `vlan_list` option was added in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure vlan router-discovery add prefix

```
configure vlan vlan_name router-discovery {ipv6} add prefix prefix
```

Description

Adds a prefix to the router discovery advertisements on the VLAN.

Syntax Description

<i>vlan_name</i>	Specifies an IPv6 configured VLAN.
<i>prefix</i>	Specifies the prefix to add.

Default

N/A.

Usage Guidelines

This command adds a prefix to the router advertisement messages for the VLAN. Prefixes defined with this command are only included in the router advertisement messages and have no operational impact on VLANs.

To configure the parameters for this prefix, use the following command:

```
configure vlan vlan_name router-discovery {ipv6} set prefix prefix
[autonomous-flag auto_on_off | onlink-flag onlink_on_off | preferred-lifetime preflife | valid-lifetime validlife]
```

Example

The following command adds the prefix 2001:db8:3456::/64 for the VLAN "top_floor":

```
configure vlan top_floor router-discovery add prefix 2001:db8:3456::/64
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure vlan router-discovery default-lifetime

```
configure vlan vlan_name router-discovery {ipv6} default-lifetime
defaultlifetime
```

Description

Configures the router lifetime value sent in router discovery advertisements on the VLAN.

Syntax Description

<i>vlan_name</i>	Specifies an IPv6 configured VLAN.
<i>defaultlifetime</i>	Specifies the router lifetime. Range is 0, or max-interval to 9000 seconds.

Default

1800 seconds.

Usage Guidelines

This command configures the router lifetime value to be included in the router advertisement messages.

The value is specified in seconds and is either 0, or between max-interval and 9000 seconds. A value of 0 indicates that the router is not to be used as a default router.

After a host sends a router solicitation, and receives a valid router advertisement with a non-zero router lifetime, the host must desist from sending additional solicitations on that interface, until an event such as re-initialization takes place.

Example

The following example configures the default-lifetime to be 3600 seconds for the VLAN "top_floor":

```
configure vlan top_floor router-discovery default-lifetime 3600
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure vlan router-discovery delete prefix

```
configure vlan vlan_name router-discovery {ipv6} delete prefix [prefix |  
all]
```

Description

Deletes prefixes from the router discovery advertisements on the VLAN.

Syntax Description

<i>vlan_name</i>	Specifies an IPv6 configured VLAN.
<i>prefix</i>	Specifies the prefix to delete.
all	Specifies to delete all prefixes.

Default

N/A.

Usage Guidelines

This command deletes previously defined router advertisement prefixes.

Example

The following example deletes the prefix 2001:db8:3161::/64 for the VLAN "top_floor":

```
configure vlan top_floor router-discovery delete 2001:db8:3161::/64
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure vlan router-discovery link-mtu

```
configure vlan vlan_name router-discovery {ipv6} link-mtu linkmtu
```

Description

Configures the link MTU value sent in router discovery advertisements on the VLAN.

Syntax Description

<i>vlan_name</i>	Specifies an IPv6 configured VLAN.
<i>linkmtu</i>	Specifies the link MTU. Range is 0 to 9216.

Default

0, meaning that no link MTU information is sent.

Usage Guidelines

This command configures the link MTU placed into the router advertisement messages. Advertisement of the MTU helps ensure use of a consistent MTU by hosts on the VLAN.

The minimum value is 0, and the maximum value is 9216. The default value is 0, which means that no link MTU information is included in the router discovery messages.

Example

The following example configures the link MTU to be 5126 for the VLAN "top_floor":

```
configure vlan top_floor router-discovery link-mtu 5126
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure vlan router-discovery managed-config-flag

```
configure vlan vlan_name router-discovery {ipv6} managed-config-flag
    on_off
```

Description

Configures the managed address configuration flag value sent in router discovery advertisements on the VLAN.

Syntax Description

<i>vlan_name</i>	Specifies an IPv6 configured VLAN.
<i>on_off</i>	Specifies setting the flag to on or off.

Default

Off.

Usage Guidelines

This command configures the contents of the managed address configuration flag in the router advertisement messages.

A value of *on* tells hosts to use the administered (stateful) protocol DHCP for address autoconfiguration in addition to any addresses autoconfigured using stateless address autoconfiguration. A value of *off* tells hosts to use stateless address autoconfiguration. If this command is not entered, the default value is off.

Example

The following example configures the managed address configuration flag to be on for the VLAN "top_floor":

```
configure vlan top_floor router-discovery managed-config-flag on
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure vlan router-discovery max-interval

```
configure vlan vlan_name router-discovery {ipv6} max-interval
maxinterval
```

Description

Configures the maximum time between unsolicited router discovery advertisements on the VLAN.

Syntax Description

<i>vlan_name</i>	Specifies an IPv6 configured VLAN.
<i>maxinterval</i>	Specifies the maximum time between advertisements, in seconds. Range is 4 to 1800.

Default

600 seconds.

Usage Guidelines

This command configures the maximum amount of time before an unsolicited router advertisement message is advertised over the links corresponding to the VLAN.

Example

The following example configures the max-interval to be 300 seconds for the VLAN "top_floor":

```
configure vlan top_floor router-discovery max-interval 300
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure vlan router-discovery min-interval

```
configure vlan vlan_name router-discovery {ipv6} min-interval
mininterval
```


Description

Configures the minimum time between unsolicited router discovery advertisements on the [VLAN](#).

Syntax Description

<i>vlan_name</i>	Specifies an IPv6 configured VLAN.
<i>mininterval</i>	Specifies the minimum time between advertisements, in seconds. Range is 3 to 1350 (see guidelines).

Default

200 seconds, or max-interval × .33 (see guidelines).

Usage Guidelines

This command configures the minimum amount of time before an unsolicited router advertisement message is advertised over the links corresponding to the VLAN.

The minimum value is three seconds. The maximum time is (.75 × max-interval) seconds. If you do not explicitly set this value, the min-interval value is reset whenever the max-interval is configured. Min-interval will then be dynamically adjusted to .33 times the max-interval.

Example

The following example configures the min-interval to be 300 seconds for the VLAN "top_floor":

```
configure vlan top_floor router-discovery min-interval 300
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure vlan router-discovery other-config-flag

```
configure vlan vlan_name router-discovery {ipv6} other-config-flag  
on_off
```

Description

Configures the other stateful configuration flag value sent in router discovery advertisements on the [VLAN](#).

Syntax Description

<i>vlan_name</i>	Specifies an IPv6 configured VLAN.
<i>on_off</i>	Specifies setting the flag to on or off.

Default

Off.

Usage Guidelines

This command configures the contents of the other stateful configuration flag in the router advertisement messages.

When set to *on*, hosts use the administered (stateful) protocol (*DHCP*) for autoconfiguration of other (non-address) information. If this command is not entered, the default value is *off*.

Example

The following example configures the other stateful configuration flag to be on for the VLAN "top_floor":

```
configure vlan top_floor router-discovery other-config-flag on
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure vlan router-discovery reachable-time

```
configure vlan vlan_name router-discovery {ipv6} reachable-time  
    reachabletime
```

Description

Configures the reachable time value in router discovery advertisements on the *VLAN*.

Syntax Description

<i>vlan_name</i>	Specifies an IPv6 configured VLAN.
<i>reachabletime</i>	Specifies the reachable time value in advertisements, in milliseconds. Range is 0 to 3,600,000 (one hour).

Default

30,000 milliseconds.

Usage Guidelines

The reachable time is the time, in milliseconds, that a node assumes a neighbor is reachable after having received a reachability confirmation. A value of 0 means the time is unspecified by this router. The maximum value is 3,600,000 (1 hour).

Example

The following example configures the reachable time to be 3,600,000 milliseconds for the VLAN "top_floor":

```
configure vlan top_floor router-discovery reachable-time 3600000
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure vlan router-discovery retransmit-time

```
configure vlan vlan_name router-discovery {ipv6} retransmit-time
retransmittime
```

Description

Configures the retransmit time value in router discovery advertisements on the VLAN.

Syntax Description

<i>vlan_name</i>	Specifies an IPv6 configured VLAN.
<i>retransmittime</i>	Specifies the reachable time value in advertisements, in milliseconds. Range is 0 to 4,294,967,295 (approximately 50 days).

Default

1,000 milliseconds.

Usage Guidelines

This command configures the retransmit time value in the router advertisement messages.

The retransmit time, in milliseconds, is the time between retransmitted neighbor solicitation messages. A value of 0 means the value is unspecified by this router. The maximum value is 4,294,967,295.

Example

The following example configures the retransmit time to be 604,800,000 milliseconds (one week) for the VLAN "top_floor":

```
configure vlan top_floor router-discovery retransmit-time 604800000
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure vlan router-discovery set prefix

```
configure vlan vlan_name router-discovery {ipv6} set prefix prefix
  [autonomous-flag auto_on_off | onlink-flag onlink_on_off | preferred-lifetime preflife | valid-lifetime validlife]
```

Description

Sets the parameters for a prefix in the router discovery advertisements on the VLAN.

Syntax Description

<i>vlan_name</i>	Specifies an IPv6 configured VLAN.
<i>prefix</i>	Specifies which prefix's parameters to set.
<i>auto_on_off</i>	Specifies the autonomous flag.
<i>onlink_on_off</i>	Specifies the on link flag.
<i>preflife</i>	Specifies the preferred lifetime in seconds. Maximum value is 4,294,967,295.
<i>validlife</i>	Specifies the valid lifetime in seconds. Maximum value is 4,294,967,295.

Default

The prefix parameter defaults are:

- Valid lifetime—2,592,000 seconds (30 days)
- On-link flag—on

- Preferred lifetime—604,800 seconds (7 days)
- Autonomous flag—on

Usage Guidelines

This command configures the attributes associated with the specified prefix.

The **autonomous-flag** flag option modifies the autonomous flag of the prefix. The autonomous flag value specifies whether the prefix can be used for autonomous address configuration (on) or not (off).

The **onlink-flag** option modifies the on link flag of the prefix. The on link flag specifies whether the prefix can be used for on link determination (on) or not (off). The default value of the on link flag is on.

The **preferred-lifetime** option modifies the preferred lifetime of a prefix. The preferred lifetime value is the time (from when the packet is sent) that addresses generated from the prefix via stateless address autoconfiguration remain preferred. The maximum value is 4,294,967,295. The default value is 604,800 seconds (7 days).

The **valid-lifetime** option modifies the valid lifetime of a prefix. The valid lifetime value is the time (from when the packet was sent) that the prefix is valid for the purpose of on-link determination. The maximum value is a 4,294,967,295. The default value is 2,592,000 seconds (30 days).

Example

The following example sets the on link parameter of the prefix 2001:db8:3161::/64 to off, for the VLAN "top_floor":

```
configure vlan top_floor router-discovery set prefix 2001:db8:3161::/64 onlink-flag off
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the in the [Switch Engine 32.3 Feature License Requirements](#) document..

configure router-discovery vrrp-lla-only

```
configure {vlan} vlan_name router-discovery {ipv6} vrrp-lla-only on_off
```

Description

Configures the router discovery advertisements to send only with VRRP link local address on the VRRP-enabled VLAN interface.

Syntax Description

<i>vlan_name</i>	Specifies an IPv6 configured VLAN.
router-discovery	IPv6 Router Discovery configuration
ipv6	IPv6 Router Discovery configuration.
vrrp-lla-only	Router advertisement is sent only with VRRP's virtual link local address.
<i>on_off</i>	Specifies setting the flag to on or off. Default is off.

Default

Default is off.

Usage Guidelines

This command configures the router advertisements to use only VRRP's link local address and avoid VLAN link local address on VRRP-enabled VLAN interfaces.

When set to on, VRRP's link local address is used in router advertisements. If this command is not entered, the default value is off and VLAN link local address is used in router advertisements.



Note

You need to explicitly set this value to "off" when VRRP is disabled on the VLAN.

Example

The following example configures the router discovery advertisements to use VRRP link local address for the VLAN "top_floor":

```
# configure vlan top_floor router-discovery vrrp-lla-only on
```

History

This command was first available in ExtremeXOS 32.3.

Platform Availability

This command is available on all platforms that support the Advanced Edge License as shown in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure vlan subvlan

```
configure vlan vlan_name [add | delete] subvlan sub_vlan_name
```

Description

Adds or deletes a subVLAN to a superVLAN.

Syntax Description

<i>vlan_name</i>	Specifies a superVLAN name.
add	Adds the subVLAN to the superVLAN.
delete	Deletes the subVLAN from the superVLAN.
<i>sub_vlan_name</i>	Specifies a subVLAN name.

Default

N/A.

Usage Guidelines

The following properties apply to VLAN aggregation operation:

- All broadcast and unknown traffic remain local to the subVLAN and does not cross the subVLAN boundary. All traffic within the subVLAN are switched by the subVLAN, allowing traffic separation between subVLANs (while using the same default router address among the subVLANs).
- Hosts can be located on the superVLAN or on subVLANs. Each host can assume any IP address within the address range of the superVLAN router interface. Hosts on the subVLAN are expected to have the same network mask as the superVLAN and have their default router set to the IP address of the superVLAN.
- All IP unicast traffic between subVLANs is routed through the superVLAN. For example, no ICMP redirects are generated for traffic between subVLANs, because the superVLAN is responsible for subVLAN routing. Unicast IP traffic across the subVLANs is facilitated by the automatic addition of an ARP entry (similar to a proxy ARP entry) when a subVLAN is added to a superVLAN. This feature can be disabled for security purposes.

Example

The following example adds the subVLAN "vsub1" to the superVLAN "vsuper":

```
configure vlan vsuper add subvlan vsub1
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure vlan subvlan-address-range

```
configure vlan vlan_name subvlan-address-range ipaddress1 ipaddress2
```

Description

Configures subVLAN address ranges on each subVLAN to prohibit the entry of IP addresses from hosts outside of the configured range.

Syntax Description

<i>vlan_name</i>	Specifies a subVLAN name.
<i>ipaddress1</i>	Specifies an IP address.
<i>ipaddress2</i>	Specifies another IP address.

Default

N/A.

Usage Guidelines

There is no error checking to prevent the configuration of overlapping subVLAN address ranges between multiple subVLANs. Doing so can result in unexpected behavior of ARP within the superVLAN and associated subVLANs.

Example

The following example configures the subVLAN vsuper to prohibit the entry of IP addresses from hosts outside of the configured range of IP addresses:

```
configure vlan vsuper subvlan-address-range 10.1.1.1 - 10.1.1.255
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure vlan suppress

```
configure vlan vlan_name suppress [arp-only | none]
```

Description

This command enables or disables ARP suppression on VXLAN tenant VLANs.

Syntax Description

<i>vlan_name</i>	VLAN name.
suppress	Specifies suppression of ARP on VXLAN tenant VLANs.
arp-only	Specifies ARP suppression. Requests may be proxied.
none	Disable ARP suppression (default).

Default

ARP is suppressed.

Usage Guidelines

This command is allowed on VXLAN tenant VLANs only.

Example

The following example enables ARP suppression on VXLAN tenant VLAN "tenant1":

```
configure vlan tenant1 suppress arp-only
```

History

This command was first available in ExtremeXOS 22.3.

Platform Availability

This command is supported on the ExtremeSwitching 5320, 5420, 5520, 5720 series switches, and stacks with 5320, 5420, 5520, 5720 slots only.

configure vlan tag

```
configure {vlan} vlan_name tag tag {remote-mirroring}
```

Description

Assigns a unique 802.1Q tag to the VLAN.

Syntax Description

<i>vlan_name</i>	Specifies a VLAN name.
<i>tag</i>	Specifies a value to use as an 802.1Q tag. The valid range is from 2 to 4095.
remote-mirroring	Specifies that the tagged VLAN is for remote mirroring.

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 use an 802.1Q tag, a tag must be assigned to the VLAN. The valid range is from 2 to 4094 (tag 1 is assigned to the default VLAN, and tag 4095 is assigned to the management VLAN).

The 802.1Q tag is also 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 becomes the VLANid for the VLAN you specify, and a new VLANid is automatically assigned to the other untagged VLAN.

Example

The following command assigns a tag (and internal VLANid) of 120 to a VLAN named accounting:

```
configure accounting tag 120
```

History

This command was first available in ExtremeXOS 10.1.

The **remote-mirroring** option was added in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure vlan udp-profile

```
configure vlan vlan_name udp-profile [profilename | none]
```

Description

Associates a UDP forwarding profile to a VLAN.

Syntax Description

<i>vlan_name</i>	Specifies a VLAN name.
<i>profilename</i>	Specifies a policy file to use for the UDP forwarding profile.
none	Removes any UDP forwarding profile from the VLAN.

Default

No UDP profiles are associated with the VLAN.

Usage Guidelines

You can apply a UDP forwarding policy only to an L3 VLAN (a VLAN having at least one IP address configured on it). If there is no IP address configured on the VLAN, then the command is rejected.

A UDP forwarding policy must contain only the following attributes. Unrecognized attributes are ignored.

- Match attributes
 - Destination UDP port number (destination-port)
 - Source IP address (source-ipaddress)
- Action modified (set) attributes
 - Destination IP address (destination-ipaddress)
 - VLAN name (vlan)

Policy files used for UDP forwarding are processed differently from standard policy files. Instead of terminating when an entry's match clause becomes true, each entry in the policy file is processed and the corresponding action is taken for each true match clause.

For example, if the following policy file is used as a UDP forwarding profile, any packets destined for UDP port 67 are sent to IP address 20.0.0.5 *and* flooded to VLAN to7:

```
entry one {
  if match all {
    destination-port 67 ;
  } then {
    destination-ipaddress 20.0.0.5 ;
  }
}
entry two {
  if match all {
    destination-port 67 ;
  } then {
    vlan "to7" ;
  }
}
```

If you include more than one VLAN set attribute or more than one destination-ipaddress set attribute in one policy entry, the last one is accepted and the rest are ignored.



Note

Although the ExtremeXOS Policy manager allows you to set a range for the destination-port, you should not specify the range for the destination-port attribute in the match clause of the policy statement for the UDP profile. If a destination-port range is configured, the last port in the range is accepted and the rest are ignored.

You can have two valid set statements in each entry of a UDP forwarding policy; one a destination-ipaddress and one a VLAN. ExtremeXOS currently allows a maximum of eight entries in a UDP

forwarding policy, so you can define a maximum of 16 destinations for one inbound broadcast UDP packet: eight IP addresses and eight VLANs.



Note

It is strongly advised to have no more than eight entries in a UDP forwarding profile. The UDP forwarding module processes those entries even if the entries do not contain any attributes for UDP forwarding. Having more than eight entries drastically reduces system performance. If the inbound UDP traffic rate is very high, having more than eight entries could cause the system to freeze or become locked.

If you rename a VLAN referred to in your UDP forwarding profile, you must manually edit the policy to reflect the new name, and refresh the policy.

You can also validate whether the UDP profile has been successfully associated with the VLAN by using the command `show policy {policy-name | detail}`. UDP forwarding is implemented as part of the netTools process, so the command does display netTools as a user of the policy.

Example

The following example associates the UDP forwarding profile "port123_to_corporate" to VLAN "to-sales":

```
configure vlan to-sales udp-profile port123_to_corporate
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms that use the Edge, Advanced Edge, or Core license. For information on the licenses available for each platform, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure vlan untagged-ports auto-move

```
configure vlan untagged-ports auto-move [on | off | inform]
```

Description

Globally, allows untagged ports to be moved directly from untagged VLANs to either different untagged VLANs or tagged VLANs.

Syntax Description

on	Auto-move global setting is <i>on</i> , which allows you to move untagged ports from untagged VLANs without first removing the port VLAN configuration.
off	Auto-move global setting is <i>off</i> ; you cannot directly move untagged ports from untagged VLANs without first removing the port VLAN configuration.
inform	Auto-move global setting is <i>on</i> , but you are informed when such a move occurs (default): Port # untagged has been auto-moved from VLAN "x" to "y".

Default

The default is **inform**.

Usage Guidelines

The global setting of this command impacts the following configuration commands:

- [configure vlan add ports](#) on page 1473
- [configure vman add ports](#) on page 1524

Moving tagged ports is not impacted by this global setting. You can move tagged ports directly without having to enable the auto-move global setting.

Example

The following example turns on the auto-move global setting:

```
configure vlan untagged-ports auto-move on
```

The following example turns on the auto-move global setting with the **inform** option:

```
configure vlan untagged-ports auto-move inform
```

When the **inform** keyword is used, you can directly move an untagged port, but you are informed that this has occurred:

```
configure vlan untagged-ports auto-move inform
configure vlan v2 add ports 1 untagged
```

Port 1 untagged has been auto-moved from VLAN "Default" to "v2".

History

This command was first available in ExtremeXOS 22.1.

The default was changed from **off** to **inform** in ExtremeXOS 22.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure vlan-translation add loopback-port

```
configure {vlan} vlan_name vlan-translation add loopback-port port
```

Description

Adds the specified port as a loopback port for the specified member VLAN.

Syntax Description

<i>vlan_name</i>	Specifies the name of the member VLAN to which you want to add the loopback port.
<i>port</i>	Specifies the port that serves as the loopback port.

Default

N/A.

Usage Guidelines

If two or more member VLANs have overlapping ports (where the same ports are assigned to both VLANs), each of the member VLANs with overlapping ports must have a dedicated loopback port.

The loopback port can be added to the member VLAN when the member VLAN is created, or you can use this command to add the loopback port at a later time.

Example

The following example adds port 2:1 as a loopback port for the member VLAN leafvlan:

```
configure leafvlan vlan-translation add loopback-port 2:1
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms that support the VLAN Translation feature. For features and the platforms that support them, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure vlan-translation add member-vlan

```
configure {vlan} vlan_name vlan-translation add member-vlan  
member_vlan_name {loopback-port port}
```

Description

Adds a member *VLAN* to a translation VLAN.

Syntax Description

<i>vlan_name</i>	Specifies the name of the translation VLAN to which you want to add the member VLAN.
<i>member_vlan_name</i>	Specifies the member VLAN to be added to the translation VLAN.
<i>port</i>	Specifies the port that serves as the loopback port.
loopback-port	If two or more member VLANs have overlapping ports (where the same ports are assigned to both VLANs), each of the member VLANs with overlapping ports must have a dedicated loopback port.

Default

N/A.

Usage Guidelines

This command configures VLAN tag translation between the two VLANs specified. The member VLAN is added to the list maintained by translation VLAN. A translation VLAN can have multiple member VLANs added to it.

Example

The following example adds member VLAN leafvlan to the translation VLAN branchvlan:

```
configure branchvlan vlan-translation add member-vlan leafvlan
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms that support the VLAN Translation feature. For features and the platforms that support them, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure vlan-translation delete loopback-port

```
configure {vlan} vlan_name vlan-translation delete loopback-port
```

Description

Deletes the loopback port from the specified member *VLAN*.

Syntax Description

<i>vlan_name</i>	Specifies the name of the member VLAN from which you want to delete the loopback port.
------------------	--

Default

N/A.

Usage Guidelines

This command disables and deletes the loopback port from the specified member VLAN. This command does not delete the member VLAN.

Example

The following example deletes the loopback port from the member VLAN leafvlan:

```
configure leafvlan vlan-translation delete loopback-port
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms that support the VLAN Translation feature. For features and the platforms that support them, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure vlan-translation delete member-vlan

```
configure {vlan} vlan_name vlan-translation delete member-vlan  
[member_vlan_name | all]
```

Description

Deletes one or all member VLANs from a translation VLAN.

Syntax Description

<i>vlan_name</i>	Specifies the name of the translation VLAN from which you want to delete the member VLAN.
<i>member_vlan_name</i>	Specifies the member VLAN to be deleted from the translation VLAN.
all	Deletes all member VLANs from the specified translation VLAN.

Default

N/A.

Usage Guidelines

This command removes the link between the translation VLAN and the specified member VLANs, but it does not remove the VLANs from the switch.

Example

The following example deletes member VLAN leafvlan from the translation VLAN branchvlan:

```
configure branchvlan vlan-translation delete member-vlan leafvlan
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms that support the VLAN Translation feature. For features and the platforms that support them, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure vm add | delete ports

```
configure vm vm_name {add | delete} ports portlist
```

Description

Adds or deletes dedicated Application Hosting (IAH) or management ports to a virtual machine (VM).

Syntax Description

vm	Designates a virtual machine.
<i>vm_name</i>	Specifies the VM name to add or delete ports to.
add	Designates adding IAH ports to a VM.
delete	Designates deleting IAH ports from a VM.
ports	Designates adding or deleting ports.
<i>portlist</i>	Selects the IAH ports to add or delete.

Default

N/A.

Usage Guidelines

Multiple VMs cannot use the same sideband port, but they can share the management port. To view ports for an existing VM, use the command `show vm {vm_name | detail}`.

This command does not take effect until the next time the guest VM is started.

The IAH feature requires the Solid State Storage Device SSD-120.

Example

The following example adds port 1-5 to VM "vm1":

```
# configure vm vm1 add ports 1-5
```

History

This command was first available in ExtremeXOS 30.3.

Platform Availability

This command is available on all platforms that support the Integrated Application Hosting (IAH) feature and have a Core or Premier license installed. For a list of platforms that support the IAH feature and for information about licenses, see the [Switch Engine 32.3 Feature License Requirements](#).

configure vm add virtual-interface

```
configure vm vm_name add virtual-interface port port {vlan vlan_id}
      {name vf_name}
```

Description

Adds a virtual interface to a guest virtual machine (VM).

Syntax Description

vm	Virtual machine.
<i>vm_name</i>	Specifies the VM to add the virtual interface to.
add	Specifies adding a virtual interface.
virtual-interface	Specifies the virtual interface to add.
name	Specifies adding an optional name to the virtual interface.
<i>vf_name</i>	Specifies an optional name (unique within this VM) for the virtual interface to add to the VM.
port	Specifies the associated dedicated port (physical function for this virtual interface).
<i>port</i>	Specifies the dedicated port number.
vlan	Specifies an optional VLAN mapped to this virtual interface.
<i>vlan_id</i>	Specifies the VLAN ID tag between 1 and 4,094.

Default

N/A.

Usage Guidelines

The maximum number of virtual interfaces that you can attach is 16. The dedicated port specified cannot already be a dedicated port within the VM.

To delete a virtual interface from a guest VM, use the `configure vm vm_name delete virtual-interface [name vf_name | mac mac_addr]` command.

Example

The following example add a virtual interface to the VM "vm1" on port 7:

```
# configure vm vm1 add virtual-interface port 7
```

History

This command was first available in ExtremeXOS 30.5.

Platform Availability

This command is available on all platforms that support the Integrated Application Hosting (IAH) feature and have a Core or Premier license installed. For a list of platforms that support the IAH feature and for information about licenses, see the [Switch Engine 32.3 Feature License Requirements](#).

configure vm delete virtual-interface

```
configure vm vm_name delete virtual-interface [name vf_name | mac mac_addr]
```

Description

Deletes a virtual interface from a guest virtual machine (VM).

Syntax Description

vm	Virtual machine.
<i>vm_name</i>	Specifies the VM to delete the virtual interface from.
delete	Specifies deleting a virtual interface.
virtual-interface	Specifies which virtual interface to delete.
name	Specifies deleting a virtual interface by specifying its optional name.
<i>vf_name</i>	Specifies the optional name (unique within this VM) of the virtual interface to delete from the VM.
mac	Specifies deleting a virtual interface by specifying its MAC address.
<i>mac_addr</i>	Specifies the virtual interface MAC address.

Default

N/A.

Usage Guidelines

To add a virtual interface to a guest VM, use the `configure vm vm_name add virtual-interface port port {vlan vlan_id} {name vf_name}` command.

Example

The following example deletes the virtual interface "my_vf" from the VM "vm1":

```
# configure vm vm1 delete virtual-interface my_vf
```

History

This command was first available in ExtremeXOS 30.5.

Platform Availability

This command is available on all platforms that support the Integrated Application Hosting (IAH) feature and have a Core or Premier license installed. For a list of platforms that support the IAH feature and for information about licenses, see the [Switch Engine 32.3 Feature License Requirements](#).

configure vm cpus

```
configure vm vm_name cpus num_cpus
```

Description

Configures an existing virtual machine (VM) CPU allocation.

Syntax Description

vm	Designates a virtual machine.
<i>vm_name</i>	Specifies the VM name to configure.
cpus	Designates specifying the number of CPUs to allocate to the VM.
<i>num_cpus</i>	Specifies the number of CPUs to allocate to the VM. Range is 1-2. The default is 1.

Default

By default, the number of CPUs allocated is 1.

Usage Guidelines

The number of CPUs allocated to a VM is set when the VM is created (default is 1), but you can change the allocation with this command. To view the number of CPUs currently allocated to a VM, use the command `show vm {vm_name | detail}`.

This command does not take effect until the next time the guest VM is started.

The Integrated Application Hosting (IAH) feature requires the Solid State Storage Device SSD-120.

Example

The following example changes the number of CPUs allocated to VM "vm1" to 2:

```
# configure vm vm1 cpus 2
```

History

This command was first available in ExtremeXOS 30.3.

Platform Availability

This command is available on all platforms that support the Integrated Application Hosting (IAH) feature and have a Core or Premier license installed. For a list of platforms that support the IAH feature and for information about licenses, see the [Switch Engine 32.3 Feature License Requirements](#).

configure vm disk bus-type

```
configure vm vm_name disk bus-type bus_type
```

Description

Configures the virtual machine's (VM's) disk bus or controller.

Syntax Description

vm	Specifies VMs.
<i>vm_name</i>	Specifies the name of the VM.
disk	Specifies disk configuration.
bus-type	Specifies configuring the disk bus type.
<i>bus_type</i>	Specifies the disk bus type (default is VirtIO).

Default

The default is VirtIO.

Usage Guidelines

This command allows you to configure the VM's disk bus or controller. The default bus type is VirtIO, but some operating systems do not support this, and as a consequence, the VM will fail to boot. In this case, you can configure the bus type to IDE or SCSI.

Example

The following example configures the VM "vm1" to the bus type to IDE:

```
# configure vm vm1 disk bus-type IDE
```

History

This command was first available in ExtremeXOS 30.7.

Platform Availability

This command is available on all platforms that support the Integrated Application Hosting (IAH) feature and have a Core or Premier license installed. For a list of platforms that support the IAH feature and for information about licenses, see the [Switch Engine 32.3 Feature License Requirements](#).

configure vm memory

```
configure vm vm_name memory memory_size
```

Description

Changes the amount of memory assigned to an existing virtual machine (VM).

Syntax Description

vm	Designates creating a virtual machine.
<i>vm_name</i>	Specifies the VM name to change memory for.
memory	Designates specifying the amount of RAM allocated to the VM.
<i>memory_size</i>	Specifies the amount of RAM (in MB) allocated to the VM. The default is 4,096.

Default

By default, the amount of RAM allocated to a VM is 4,096.

Usage Guidelines

The amount of RAM allocated to a VM is set when the VM is created (default is 4,096 MB), but you can change the allocation with this command. To view the amount of RAM currently allocated to a VM, use the command `show vm {vm_name | detail}`.

This command does not take effect until the next time the guest VM is started.

The Integrated Application Hosting (IAH) feature requires the Solid State Storage Device SSD-120.

Example

The following example changes the amount of RAM allocated to VM "vm1" to 2,000 MB:

```
# configure vm vm1 memory 2000
```

History

This command was first available in ExtremeXOS 30.3.

Platform Availability

This command is available on all platforms that support the Integrated Application Hosting (IAH) feature and have a Core or Premier license installed. For a list of platforms that support the IAH feature and for information about licenses, see the [Switch Engine 32.3 Feature License Requirements](#).

configure vm vnc

```
configure vm vm_name vnc [none | vnc_display]
```

Description

Configures the VNC display for a virtual machine (VM).

Syntax Description

vm	Designates a virtual machine.
<i>vm_name</i>	Specifies the VM name to configure the VNC display for.
vnc	Specifies providing a display number for VNC access.
none	Disables VNC access (default).
<i>vnc_display</i>	Specifies the VNC screen number. Range is 0–15.

Default

By default, VNC access is disabled.

Usage Guidelines

For the VNC display number (or screen number), you can use the values from 0 to 15. These correspond to TCP ports 5,900 to 5,915.

Multiple VMs can be configured with the same VNC display, but VMs configured with the same display number cannot run at the same time. A VM cannot be started if the VNC port is already in use.

For security reasons, the VNC display is only accessible using SSH tunnel.

Example

The following example enables VNC on VM "vm1" with display number 3:

```
# configure vm vm1 vnc 3
```

History

This command was first available in ExtremeXOS 30.5.

Platform Availability

This command is available on all platforms that support the Integrated Application Hosting (IAH) feature and have a Core or Premier license installed. For a list of platforms that support the IAH feature and for information about licenses, see the [Switch Engine 32.3 Feature License Requirements](#).

configure vman add ports

```
configure vman vman_name add ports [port_list | all] {tagged | untagged
    {port-cvid port_cvid} | cep [ cvid cvid_first { - cvid_last }
    { translate cvid_first_xlate { - cvid_last_xlate } } | port-cvid
    port_cvid ]
```

Description

Adds one or more ports to a VMAN.

Syntax Description

<i>vman-name</i>	Specifies the name of the VMAN to configure.
<i>vman_id</i>	Specifies the ID of the VMAN to configure
all	Specifies all switch ports.
<i>port_list</i>	Specifies a list of ports.
untagged	Configures the specified ports as Customer Network Ports (CNPs).
tagged	Configures the specified ports as Provider Network Ports (PNPs), which are also called VMAN network ports.
port-cvid	Port's customer <u>VLAN</u> ID used for untagged packets.
<i>port_cvid</i>	Customer VLAN ID assigned to untagged packets from 1.

Default

If you do not specify a parameter, the default value is untagged, which creates a CNP.

Usage Guidelines

This command adds ports as either CNPs or PNPs. To add a port to a VMAN as a CEP, use the following command:

[configure vman add ports cep](#) on page 1526

The VMAN must already exist before you can add (or delete) ports. VMAN ports can belong to load-sharing groups.

When a port is configured serve as a CNP for one VMAN and A PNP for another VMAN, it inspects the VMAN ethertype in received packets. Packets with a matching ethertype are treated as tagged and switched across the associated PNP VMAN. Packets with a non-matching ethertype are treated as untagged and forwarded into the associated CNP VMAN.

When a port is configured only as a CNP (an untagged VMAN member), whether the VMAN ethertype is 0x8100 or otherwise, all received packets ingress the associated VMAN regardless of the packet's tagging.



Note

If you use the same name across categories (for example, [STPD](#) and [EAPS](#) names), we recommend that you specify the identifying keyword as well as the actual name. If you do not use the keyword, the system may return an error message.

The following guidelines apply to all platforms:

- You must enable or disable jumbo frames before configuring VMANs. You can enable or disable jumbo frames on individual ports or on the entire switch. See “Configuring Ports on a Switch” in the [Switch Engine 32.3 User Guide](#) for more information on configuring jumbo frames.
- Each port can serve in only one VMAN role per VMAN. When multiple roles are configured on a port, each role must be configured for a different VMAN.
- Multiple VMAN roles can be combined on one port with certain VLAN types as shown in the following table.

Example

The following example assigns ports 1:1, 1:2, 1:3, and 1:6 to a VMAN named accounting:

```
configure vman accounting add ports 1:1, 1:2, 1:3, 1:6 tag 100
```

History

This command was first available in ExtremeXOS 11.0.

The **cvid** keyword was added in ExtremeXOS 15.3.2.

The *vman_id* variable was added in ExtremeXOS 16.1.

The **cvid** keyword was removed in ExtremeXOS 21.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure vman add ports cep

```
configure [{vman} vman_name | vman vman_id] add ports port_list
    cep cvid cvid_first {- cvid_last} {translate cvid_first_xlate {-
    cvid_last_xlate }} | port-cvid port_cvid ]}
```

Description

Adds one or more switch ports to the specified VMAN as Customer Edge Ports (CEPs), and configures the CVIDs on those ports to map to the VMAN.

Syntax Description

<i>vman_name</i>	Specifies the VMAN to configure.
<i>vman_id</i>	Specifies the VMAN ID to configure.
<i>port_list</i>	Specifies a list of ports.
<i>cvid_first</i>	Specifies a CVLAN ID (CVID) or the first in a range of CVIDs that the CEP will accept and map to the specified VMAN. Valid values are 1-4095.
<i>cvid_last</i>	Specifies the last in a range of CVIDs that the CEP will accept and map to the VMAN. Valid values are 1-4095.
translate	Enables translation of the specified CEP CVID range to the specified VMAN CVID range.
<i>cvid_first_xlate</i>	Specifies a VMAN CVID or the first in a range of VMAN CVIDs to which the CEP CVIDs will map. Valid values are 1-4095.
<i>cvid_last_xlate</i>	Specifies the last in a range of VMAN CVIDs to which the CEP CVIDs will map. Valid values are 1-4095. The number of VMAN CVIDs in this range must equal the number of CEP CVIDs specified in this command.
port-cvid	Port's customer <u>VLAN</u> ID used for untagged packets.
<i>port_cvid</i>	Customer VLAN ID assigned to untagged packets from 1.

Default

N/A.

Usage Guidelines

If you specify only one CVID or a range of CVIDs without translation, the specified CVIDs are mapped to the specified VMAN and appear unchanged in the VMAN.

If you specify CVID translation, the CEP CVIDs map to different VMAN CVIDs. The number of CEP CVIDs specified must equal the number of VMAN CVIDs specified. The first CEP CVID in the specified

range maps to the first CVID in the range specified for the VMAN. The difference between `cvid_first` and `cvid_first_xlate` establishes an offset N that maps CEP CVIDs to VMAN CVIDs. (Offset N = `cvid_first_xlate` - `cvid_first`.) The translated VMAN CVID that corresponds to a CEP CVID can be determined as follows:

VMAN CVID = CEP CVID + N

**Note**

CVID translation can reduce the number of CVIDs that can be mapped to VMANs.

After you enable and configure a CEP with this command, you can use the following command to map additional CVIDs on the port to the VMAN:

```
configure [ {vman} vman_name | vman vman_id] ports port_list
add cvid cvid_first {- cvid_last} {translate cvid_first_xlate {-
cvid_last_xlate}}
```

When this command specifies multiple ports, each port gets an independent CVID map; the ports do not share a common map. Changes to the CVID map affect only the ports specified in the configuration command. For example, consider the following commands:

```
configure vman vman1 add port 1-2 cep cvid 10
```

```
configure vman vman1 port 1 add cvid 11
```

After these commands are entered, port 1 maps CVIDs 10 and 11 to VMAN `vman1`, and port 2 maps only CVID 10 to `vman1`.

You can add the same port as a CEP to multiple VMANs. A port can also support multiple VMANs in different roles as shown in `configure vman vman_name add ports`.

To view the CEP CVID configuration for a port, use the `show vman` command.

ExtremeXOS 21.1 adds an optional port CVID parameter to the existing untagged and CEP VMAN port configuration options. When present, any untagged packet received on the port will be double tagged with the configured port CVID and the SVID associated with the VMAN. If the port is untagged, packets received with a single CVID will still have the SVID added as usual. If the port is CEP, only untagged and any specifically configured CVIDs will be allowed. As double tagged packets are received from tagged VMAN ports and forwarded to untagged VMAN ports, the SVID associated with the VMAN is stripped. Additionally, the CVID associated with the configured Port CVID is also stripped in the same operation. If the port is CEP and CEP egress filtering is enabled, only the specified port-cvid and cvids are allowed to egress.

Example

The following example configures port 1 as a CEP for VMAN `vman1` and specifies that CEP CVID 5 maps to CVID 5 on the VMAN:

```
configure vman vman1 add port 1 cep cvid 5
```

The following example configures port 1 as a CEP for VMAN vman1 and enables the port to translate CEP CVIDs 10-19 to VMAN CVIDs 20-29:

```
configure vman vman1 add port 1 cep cvid 10 - 19 translate 20 - 29
```

History

This command was first available in ExtremeXOS 12.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

The CVID translation feature is available on all platforms.

configure vman delete ports

```
configure vman [vman_name | vman_list] delete ports [all | port_list]
```

Description

Deletes one or more ports from a VMAN.

Syntax Description

<i>vman_name</i>	Specifies a VMAN name.
<i>vman_list</i>	Specifies a VMAN list name.
all	Specifies all ports in the VMAN.
<i>port_list</i>	Specifies a list of ports.

Default

N/A.

Usage Guidelines

The VMAN must already exist before you can delete ports.

Example

The following example deletes ports 1, 2, 3, and 6 on a switch for a VMAN named accounting:

```
configure vman accounting delete ports 1,2,3,6
```

History

This command was first available in ExtremeXOS 11.0.

The `vman_list` variable was added in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure vman ethertype

```
configure vman ethertype value [primary | secondary]
```

Description

Changes the default ethertype for the VMAN header.

Syntax Description

<i>value</i>	Specifies an ethertype value in the format of 0xffff.
primary	Assigns the ethertype as the primary Ethernet value.
secondary	Assigns the ethertype as the secondary Ethernet value.

Default

Ethertype value of 0x88a8 and type primary.

Usage Guidelines

The software supports two VMAN ethertype values: a primary value and a secondary value. By default, the primary ethertype applies to all VMANs. To use the secondary ethertype, define the ethertype with this command, and then assign the secondary ethertype to ports with the following command:

```
configure port port_list ethertype {primary | secondary}
```

If your VMAN transits a third-party device (other than an Extreme Networks device), you must configure the ethertype for the VMAN tag as the ethertype that the third-party device uses. If you configure both primary and secondary ethertypes, you can connect to devices that use either of the two values assigned.

The system supports all VMAN ethertypes, including the standard ethertype of 0x8100.

Example

The following command changes the VMAN ethertype value to 8100:

```
configure vman ethertype 0x8100
```

History

This command was first available in ExtremeXOS 11.0.

Support for a secondary ethertype was added in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure vman ports add cvid

```
configure vman vman_name ports [port_list | all] add [cvid cvid_first
    { - cvid_last } { translate cvid_first_xlate { - cvid_last_xlate } }
    | port-cvid port_cvid]
```

Description

Adds one or more CVIDs to a CEP.

Syntax Description

<i>vman_name</i>	Specifies the VMAN to configure.
<i>vman_id</i>	Specifies the VMAN ID to configure.
<i>port_list</i>	Specifies a list of ports.
<i>cvid_first</i>	Specifies a Customer VLAN ID (CVID) or the first in a range of CVIDs that the CEP will accept and map to the specified VMAN. Valid values are 1-4095.
<i>cvid_last</i>	Specifies the last in a range of CVIDs that the CEP will accept and map to the VMAN. Valid values are 1-4095.
translate	Enables translation of the specified CEP CVID range to the specified VMAN CVID range.
<i>cvid_first_xlate</i>	Specifies a VMAN CVID or the first in a range of VMAN CVIDs to which the CEP CVIDs will map. Valid values are 1-4095.
<i>cvid_last_xlate</i>	Specifies the last in a range of VMAN CVIDs to which the CEP CVIDs will map. Valid values are 1-4095. The number of VMAN CVIDs in this range must equal the number of CEP CVIDs specified in this command.
port-cvid	Port's customer VLAN ID used for untagged packets.
<i>port_cvid</i>	Customer VLAN ID assigned to untagged packets from 1.

Default

N/A.

Usage Guidelines

Before you can add CVIDs to CEPs, you must configure the target physical ports as CEPs using the following command:

[configure vman add ports](#) on page 1524

If you specify only one CVID or a range of CVIDs without translation, the specified CVIDs are mapped to the specified VMAN and appear unchanged in the VMAN.

If you specify CVID translation, the CEP CVIDs map to different VMAN CVIDs. The number of CEP CVIDs specified must equal the number of VMAN CVIDs specified. The first CEP CVID in the specified range maps to the first CVID in the range specified for the VMAN. The difference between `cvid_first` and `cvid_first_xlate` establishes an offset *N* that maps CEP CVIDs to VMAN CVIDs. (Offset *N* = `cvid_first_xlate` - `cvid_first`.) The translated VMAN CVID that corresponds to a CEP CVID can be determined as follows:

$$\text{VMAN CVID} = \text{CEP CVID} + N$$



Note

CVID translation can reduce the number of CVIDs that can be mapped to VMANs.

When this command specifies multiple ports, each port gets an independent CVID map; the ports do not share a common map. Changes to the CVID map affect only the ports specified in the configuration command. For example, consider the following commands:

```
configure vman vman1 add port 1-2 cep cvid 10
configure vman vman1 port 1 add cvid 11
```

After these commands are entered, port 1 maps CVIDs 10 and 11 to VMAN `vman1`, and port 2 maps only CVID 10 to `vman1`.

To view the CEP CVID configuration for a port, use the `show vman` command.

Example

The following example adds CVIDs 20-29 to port 1 and VMAN `vman1` and enables translation to CVIDs 30-39:

```
configure vman vman1 port 1 add cvid 20 - 29 translate 30 - 39
```

History

This command was first available in ExtremeXOS 12.6.

The `vman_id` variable was added in ExtremeXOS 16.1.

Platform Availability

This command is available on all platform.

configure vman ports delete cvid

```
configure vman vman_name ports [port_list | all] delete [cvid cvid_first
{ - cvid_last } | port-cvid port_cvid]
```

Description

Deletes one or more CVIDs from a CEP.

Syntax Description

<i>vman_name</i>	Specifies the VMAN to configure.
<i>vman_list</i>	Specifies the VMAN list to configure.
<i>port_list</i>	Specifies a list of ports.
<i>cvid_first</i>	Specifies a CVID or the first in a range of CVIDs that are to be deleted. Valid values are 1-4095.
<i>cvid_last</i>	Specifies the last in a range of CVIDs that are to be deleted. Valid values are 1-4095.

Default

N/A.

Usage Guidelines

Each CEP has its own CVID map, and this command deletes CVIDs only from the ports specified with this command.

If all the CVIDs are deleted from a CEP, the CEP is deleted from the VMAN.

To view the CEP CVID configuration for a port, use the `show vman` command.

Example

The following command deletes CVID 15 on port 1 from VMAN vman1:

```
configure vman vman1 port 1 delete cvid 15
```

History

This command was first available in ExtremeXOS 12.6.

The *vman_list* variable was added in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure vman protocol

```
configure vman [vman_name | vman_list] protocol {filter} filter_name
```


Description

Configures a VMAN to use a specific protocol filter.

Syntax Description

<i>vman_name</i>	Specifies a VMAN name.
<i>vman_list</i>	Specifies a VMAN list.
protocol	Specifies a protocol filter.
filter	Specifies a protocol filter.
<i>filter_name</i>	Specifies a protocol filter name.

Default

N/A.

Usage Guidelines

Use this command to configure a VMAN to use a specific protocol filter.

Protocol Filters

These devices do not forward packets with a protocol-based *VLAN* set to AppleTalk. To ensure that AppleTalk packets are forwarded on the device, create a protocol-based VLAN set to “any” and define other protocol-based VLANs for other traffic, such as IP traffic. The AppleTalk packets pass on the “any” VLAN, and the other protocols pass traffic on their specific protocol-based VLANs.

Example

The following example configures the protocol filter “my_filter” to vlan v1:

```
configure vlan v1 protocol my_filter
configure vlan v1 protocol filter my_filter
```

History

This command was first available in ExtremeXOS 10.1.

The **filter** keyword was added in ExtremeXOS 15.5.

The *vman_list* variable was added in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure vman tag

```
configure vman vman_name tag tag
```

Description

Assigns a tag to a VMAN.

Syntax Description

<i>vman_name</i>	Specifies a VMAN name.
<i>tag</i>	Specifies a value to use as the VMAN tag. The valid range is from 2 to 4094.

Default

N/A.

Usage Guidelines

Every VMAN requires a unique tag.

You can specify a value that is currently used as an internal *VLAN* ID on another VLAN; it becomes the VLAN ID for the VLAN you specify, and a new VLAN ID is automatically assigned to the other untagged VLAN.

Example

The following example assigns a tag of 120 to a VMAN named "accounting":

```
configure vman accounting tag 120
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure vm-tracking authentication database-order

```
configure vm-tracking authentication database-order [[nms] | [vm-map] |  
[local] | [nms local] | [local nms] | [nms vm-map] | [vm-maplocal] |  
[local vm-map] | [nms vm-map local] | [localnmsvm-map]]
```

Description

Configures the authentication database options and sequence for VM authentication.

Syntax Description

nms	Specifies the configured Network Management System (NMS).
vm-map	Specifies the configured VMMAP file.
local	Specifies the configured local database.

Default

nms vm-map local.

Usage Guidelines

The switch attempts VM authentication in the sequence specified. For example, in the default configuration, the switch attempts NMS authentication first, VMMAP authentication second, and local authentication third. If nms is specified, the switch always attempts NMS authentication before attempting VMMAP file authentication.

Example

The following command configures the database authentication order:

```
# configure vm-tracking authentication database-order local nms vm-map
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure vm-tracking blackhole

```
configure vm-tracking blackhole [policy policy_name | dynamic-rule  
                                rule_name | none]
```

Description

Specifies a policy file or dynamic *ACL* rule to apply to VMs during periods that are outside of the approved time slot for that VM.

Syntax Description

<i>policy_name</i>	Specifies the name of a policy file to apply to the VM authentication request.
<i>rule_name</i>	Specifies the name of an ACL rule to apply to the VM authentication request.

Default

N/A.

Usage Guidelines

This command is not supported in this software release. It will be supported in a future release.

The none option applies no policy name or ACL rule during periods that are outside of the approved time slot for that VM.



Note

This command is provided to support future identity management features. It serves no practical purpose in this release.

Example

The following command applies no policy name or ACL rule during periods that are outside of the authorized authentication period:

```
# configure vm-tracking blackhole none
```

History

This command was first visible in ExtremeXOS 12.5, but it is not supported in this release. This command will be supported in a future release.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure vm-tracking local-vm

```
configure vm-tracking local-vm mac-address mac [name name | ip-address ipaddress | vpp vpp_name] | vlan-tag tag {vr vr_name}}
```

Description

Configures the parameters associated with a local VM database entry to be used for VM MAC local authentication.

Syntax Description

<i>mac</i>	Specifies the MAC address for the VM database entry you want to configure.
<i>name</i>	Specifies a name to represent this VM in show vm-tracking command display.
<i>ipaddress</i>	Specifies the IP address for the VM. This must match the IP address configured on the VM.
<i>vpp_name</i>	Specifies the name of a VPP to apply to the local VM.
<i>tag</i>	<u>VLAN</u> tag between 1 and 4094.
<i>vr_name</i>	Virtual router name.

Default

N/A.

Usage Guidelines

Before you configure a VM entry in the local VM database, you must create the entry with the `create vm-tracking local-vm` command.

Before you assign an VPP to a VM entry in the local VM database, you must create the VPP with the `create vm-tracking vpp` command.

Example

The following command configures an IP address for the VM entry specified by the MAC address:

```
# configure vm-tracking local-vm mac-address 00:E0:2B:12:34:56 ip-address 10.10.10.1
```

History

This command was first available in ExtremeXOS 12.5.

The **ingress-vpp** and **egress-vpp** options were replaced with the vpp option in ExtremeXOS 12.6.

The **vlan-tag** and **vr-name** options were added in ExtremeXOS15.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure vm-tracking nms timeout

```
configure vm-tracking nms timeout seconds
```

Description

Configures the timeout period for authentication attempts with the configured NMS servers.

Syntax Description

<i>seconds</i>	Specifies the timeout period in seconds.
----------------	--

Default

3 seconds.

Usage Guidelines

None.

Example

The following command configures the switch to allow 1 minute for successful authentication of a VM with the NMS server:

```
# configure vm-tracking nms timeout 60
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure vm-tracking nms

```
configure vm-tracking nms [primary | secondary] server [ipaddress |  
  hostname] {udp_port} client-ip client_ip shared-secret {encrypted  
  encrypted_secret | secret } {vr vr_name}
```

Description

Configures the switch *RADIUS* client to an NMS for VM authentication.

Syntax Description

primary secondary	Specifies the whether you are configuring the primary or secondary NMS.
<i>ipaddress</i>	Specifies the NMS IP address.
<i>hostname</i>	Specifies the NMS DNS hostname.

<i>udp_port</i>	Specifies the UDP port number of the NMS application.
<i>client_ip</i>	Specifies the client IP address, which is the switch IP address on the interface leading to the NMS.
encrypted	Specifies that the secret key for communications with the NMS is encrypted.
<i>secret</i>	Specifies a key or password for communications with the NMS.
<i>vr_name</i>	Specifies the VR that is used to access the NMS.

Default

N/A.

Usage Guidelines

The NMS is a RADIUS server such as the one provided with Ridgeline.

Example

The following command configures the switch to authenticate VMs through the primary NMS server Ridgeline using the password password:

```
# configure vm-tracking nms primary server Ridgeline client-ip 10.10.3.3 shared-secret password
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure vm-tracking repository

```
configure vm-tracking repository [primary | secondary] server [ipaddress
| hostname] {vr vr_name} {refresh-interval seconds} {path-name
path_name} {user user_name {encrypted encrypted_password | password }
```

Description

Configures FTP file synchronization for NVPP and VMMAF files.

Syntax Description

primary secondary	Specifies the whether you are configuring the primary or secondary FTP server.
<i>ipaddress</i>	Specifies the FTP server IP address.

<i>vr_name</i>	Specifies the VR that is used to access the FTP server.
<i>seconds</i>	Specifies how often the switch updates the local files that are synchronized with the FTP server. The range is 40 to 3600 seconds.
<i>path_name</i>	Specifies the path to the repository server files from the FTP server root directory. The default directory for repository server files is: pub.
<i>user_name</i>	Specifies a user name for FTP server access. If no username is specified, the switch uses user name anonymous.
encrypted	This keyword indicates that the specified password is encrypted.
<i>password</i>	Specifies the password for the specified user name.

Default

Refresh interval: 600 seconds.

Usage Guidelines

Some jitter is added to the refresh interval period to prevent all switches from downloading files at the same time.

Example

The following example configures the switch to refresh the VM MAP and NVPP files from primary FTP server ftp1 every five minutes:

```
# configure vm-tracking repository primary server ftp1 refresh-interval 300
```

History

This command was first available in ExtremeXOS 12.5.

Support for specifying an FTP user name was added in ExtremeXOS 12.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure vm-tracking timers

```
configure vm-tracking timers reauth-period reauth_period
```

Description

Configures the RADIUS reauthentication period for VM MAC addresses.

Syntax Description

<i>reauth_period</i>	Specifies the reauthentication period in seconds. The ranges are 0 and 30-7200 seconds.
----------------------	---

Default

0 seconds.

Usage Guidelines

One way to periodically apply Virtual Port Profiles (VPPs) to VM MAC addresses is to configure a reauthentication period. At the end of each reauthentication period, the switch reauthenticates each VM MAC address and applies any updated VPPs.

This command applies to only those VMs that authenticate through RADIUS. Reauthentication is disabled when the reauthentication period is set to 0 seconds. When reauthentication is disabled, the VM MAC address remains authenticated until the *FDB* entry for that VM expires.

Example

The following command enables RADIUS server reauthentication at 2 minute intervals:

```
# configure vm-tracking timers reauth-period 120
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure vm-tracking vpp add

```
configure vm-tracking vpp vpp_name add [ingress | egress] [policy
policy_name | dynamic-rule rule_name] {policy-order policy_order}
```

Description

Configures an LVPP to use the specified policy or *ACL* rule.

Syntax Description

<i>vpp_name</i>	Specifies the name of an existing LVPP.
add	Specifies whether the LVPP should start using the specified policy or rule.
ingress	Specifies that the policy mapped to the LVPP is for ingress traffic.

egress	Specifies that the policy mapped to the LVPP is for egress traffic.
<i>policy_name</i>	Specifies a policy to add to or delete from the LVPP.
<i>rule_name</i>	Specifies a dynamic ACL rule to add to or delete from the LVPP.

Default

N/A.

Usage Guidelines

Multiple ACL or policy files can be mapped to each LVPP. A maximum of 8 ingress and 4 egress ACL or policies are available to be mapped to each LVPP. If the policy file or dynamic rule specified in this command fails to bind, then the CLI command is rejected.

Before you can configure an LVPP, you must first create it with the `create vm-tracking vpp` command.

Example

The following command configures LVPP vpp1 to use the dynamic ACL rule named rule1 for ingress traffic:

```
# configure vm-tracking vpp vpp1 add ingress dynamic-rule rule1
```

History

This command was first available in ExtremeXOS 12.5.

The ingress and egress keywords were added in ExtremeXOS 12.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure vm-tracking vpp counters

```
configure vm-tracking vpp vpp_name counters [ingress-only | egress-only  
| both | none]
```

Description

Configures whether counters need to be installed for Virtual Machine MAC which receives this VPP mapping.

Syntax Description

ingress-only	Only counts packets ingressing the switch whose source MAC address matches VM MAC.
egress-only	Only counts packets egressing the switch whose source MAC address matches VM MAC.
both	Counts packets ingressing and egressing the switch whose source MAC address matches VM MAC.
none	No packets will be counted.

Default

N/A.

Usage Guidelines

Use this command to configure whether counters need to be installed for Virtual Machine MAC which receives this VPP mapping.

Example

Example output not yet available and will be provided in a future release.

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure vm-tracking vpp delete

```
configure vm-tracking vpp vpp_name delete [ingress | egress] [policy
policy_name | dynamic-rule rule_name] {policy-order policy_order}
```

Description

Specifies that the LVPP should stop using the specified policy or rule.

Syntax Description

<i>vpp_name</i>	Specifies the name of an existing LVPP.
delete	Specifies whether the LVPP should stop using the specified policy or rule.
ingress	Specifies that the policy mapped to the LVPP is for ingress traffic.

egress	Specifies that the policy mapped to the LVPP is for egress traffic.
<i>policy_name</i>	Specifies a policy to add to or delete from the LVPP.
<i>rule_name</i>	Specifies a dynamic ACL rule to add to or delete from the LVPP.

Default

N/A.

Usage Guidelines

Multiple ACL or policy files can be mapped to each LVPP. A maximum of 8 ingress and 4 egress ACL or policies are available to be mapped to each LVPP. If the policy file or dynamic rule specified in this command fails to bind, then the CLI command is rejected.

Before you can configure an LVPP, you must first create it with the `create vm-tracking vpp` command.

Example

The following command configures LVPP vpp1 to use the dynamic ACL rule named rule1 for ingress traffic:

```
# configure vm-tracking vpp vpp1 add ingress dynamic-rule rule1
```

History

This command was first available in ExtremeXOS 12.5.

The ingress and egress keywords were added in ExtremeXOS 12.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure vm-tracking vpp vlan-tag

```
configure vm-tracking vpp vpp_name vlan-tag tag {vr vr_name}
```

Description

This command configures the [VLAN](#) tag and VR name for VPP. If the detected VM MAC uses this VPP, then the port in which the VM MAC is detected will be placed on this VR/VLAN.

Syntax Description

<i>vpp_name</i>	Specifies a name for the LVPP.
<i>tag</i>	Specifies a name for the VLAN tag.
<i>vr_name</i>	Specifies a name for the Virtual Router.

Default

N/A.

Usage Guidelines

Use this command to configure the VLAN tag and VR name for VPP. If the detected VM MAC uses this VPP, then the port in which the VM MAC is detected will be placed on this VR/VLAN.

Example

Example output not yet available and will be provided in a future release.

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure vpex auto-configuration mlag-id

```
configure vpex auto-configuration mlag-id [cascade | ring]
```

Description

Sets a preference for when Auto-Configuration configures the virtual MLAG ID. The default (cascade) configures the virtual MLAG ID whenever a 2nd-tier BPE is detected remotely, but not locally.

Syntax Description

vpex	Specifies Virtual Port Extender (VPEX).
auto-configuration	Specifies the VPEX Auto-Configuration of cascade ports and VPEX slots.
mlag-id	Specifies setting a unique MLAG identifier of the MLAG port attached to the bridge port extender (BPE).

cascade	Specifies to configure the virtual MLAG identifier when a 2nd-tier BPE is detected on the other controlling bridge, but not locally (default).
ring	Specifies to configure the virtual MLAG identifier when a ring is detected.

Default

cascade.

Usage Guidelines

The **ring** option provides more cabling time if you need to cable the chain of BPEs as duel-homed instead of single-homed.

This command is applicable only when VPEX Auto-Configuration is enabled in MLAG mode.

Example

The following example configures the default cascade option:

```
# configure vpex auto-configuration mlag-id cascade
```

The following example displays output when VPEX Auto-Configuration is not enabled in MLAG mode:

```
# configure vpex auto-configuration mlag-id ring
Error: VPEX Auto-Configuration is not in MLAG mode. Bring up MLAG peer, disable then
re-enable VPEX Auto-Configuration.
```

History

This command was first available in ExtremeXOS 31.7.

Platform Availability

This command is available on ExtremeSwitching 5420 and 5520 series switches.

configure vpex mlag-id peer

```
configure vpex mlag-id mlag_id peer peer_name slot slot_num
```

Description

In an Extended Edge Switching topology, allows the bridge port extender (BPE) slot assignment to be applied to an identifier on the specified MLAG peer when the port connected to the BPE is physically connected to the MLAG peer switch.

Syntax Description

vpex	Specifies Virtual Port Extender (VPEX).
mlag-id	Specifies setting a unique MLAG identifier of the MLAG port attached to the bridge port extender (BPE).
<i>mlag_id</i>	Sets the MLAG identifier value of the MLAG port attached to the BPE. Range is 1-65,000.
peer	Specifies naming the MLAG peer switch.
<i>peer_name</i>	Name of the MLAG peer switch.
slot	Specifies configuring the slot identifier for the attached BPE.
<i>slot_num</i>	Specifies the BPE slot number. Range is 100-162.

Default

N/A

Usage Guidelines

An Extended Edge Switching topology allows the BPE slot assignment to be applied to an MLAG identifier on the specified MLAG peer when the port connected to the BPE is physically connected to the MLAG peer switch.

The same Extended Edge Switching slot number must have been declared on the MLAG peer that has a port in the MLAG. On the peer with the MLAG port, either this form of the command can be used, or the traditional form where a controlling bridge port is related to a slot number.

Example

The following example for MLAG peer switch "cb2" declares slot 100 on MLAG "11":

```
# configure vpex mlag-id 11 peer cb2 slot 100
```

History

This command was first available in ExtremeXOS 22.7.

Platform Availability

This command is available on ExtremeSwitching 5420 and 5520 series switches.

configure vpex ports

```
configure vpex ports port_list slot slot_num
```

Description

Allows you to associate a bridge port extender (BPE) to a slot.

Syntax Description

vpex	Specifies Virtual Port Extender (VPEX).
ports	Specifies the switch ports attached to the BPE.
<i>port_list</i>	Specifies the switch ports attached to the BPE. Must be in the format <i>slot:port</i> . Only a single port can be configured at a time. Note: If the switch port is a LAG, the port specified must be the master port.
slot	Specifies VPEX BPE slot assignment.
<i>slot_num</i>	Specifies VPEX BPE slot assignment. Value must be between 100-162.

Default

N/A

Usage Guidelines

You must enable VPEX mode (`enable vpex`) before using this command.

The behavior of this command is similar to assigning slots within a chassis. After assigning a slot number to the port extender, you can make port-level configuration choices with the familiar *slot:port* notation in other commands involving the port extender's ports (for example, `configure vlan v1 add port 100:1`).

This command causes jumbo frames to be enabled on the specified ports.

Example

The following example assigns a BPE attached to switch port 1:23 to slot 100:

```
# configure vpex ports 1:23 slot 100
```

History

This command was first available in ExtremeXOS 22.5.

Platform Availability

This command is available on ExtremeSwitching 5420 and 5520 series switches.

configure vpex ring rebalancing

```
configure vpex ring rebalancing [auto | off]
```


Description

In an Extended Edge Switching ring topology, places the "ring common" link between approximately equal length cascades.

Syntax Description

vpex	Specifies Virtual Port Extender (VPEX).
ring	Specifies ring topology changes.
rebalancing	Places the ring common link between approximately equal length cascades.
auto	Ring re-balancing will automatically run at the next ring convergence.
off	Automatic ring re-balancing is disabled (default).

Default

By default, ring re-balancing is disabled.

Usage Guidelines

This command controls the Extended Edge Switching ring re-balancing operation. Re-balancing may or may not take place at the time that ring formation is complete, depending on the setting of this command. An Extended Edge Switching ring consists of two configured Extended Edge Switching cascades of BPEs that are connected at their ends. The connected link is called the ring common link. The ring forms automatically. Two control plane cascades span all bridge port extenders (BPEs) in the ring, with each originating from a controlling bridge (CB) port and ending at the BPE that is connected to the other CB port. However, the data plane cascades remain as configured (that is, no data plane traffic crosses the common link). Re-balancing moves the ring common link so that the data plane cascades are approximately equal in length. The cost of doing this is a data plane disruption to some BPEs in the ring that is the same as that which would have occurred had a single link in the ring been broken. Re-balancing is a dynamic operation. It does not change the cascade configurations.

Changing this setting takes effect the next time that a ring experiences a new ring formation. There is no immediate effect.

You can view your re-balancing selection with the `show vpex` command.

Example

The following example turns off ring re-balancing:

```
# configure vpex ring rebalancing off
```

History

This command was first available in ExtremeXOS 22.7.

Platform Availability

This command is available on ExtremeSwitching 5420 and 5520 series switches.

configure vpls

```
configure vpls vpls_name {dot1q [ethertype hex_number | tag [include | exclude]]} {mtu number}
```



Note

This command has been replaced with the following command: `configure l2vpn [vpls vpls_name | vpwsvpws vpws_name] {dot1q [ethertype hex_number | tag [include | exclude]]} {mtu number}`.

This command is still supported for backward compatibility, but it will be removed from a future release, so we recommend that you start using the new command.

Description

Configures VPLS parameters.

Syntax Description

<i>vpls_name</i>	Identifies the VPLS within the switch (character string).
dot1q	Specifies the action the switch performs with respect to the 802.1Q ethertype or tag.
ethertype	Overwrites the ethertype value for the customer traffic sent across the PW.
<i>hex_number</i>	Identifies the ethertype, uses the format of 0xN.
tag	Specifies the action the switch performs with respect to the 802.1Q tag.
include	Includes the 802.1Q tag when sending packets over the VPLS L2 VPN.
exclude	Strips the 802.1Q tag before sending packets over the VPLS L2 VPN.
mtu	Specifies the MTU value of the VPLS transport payload packet.
<i>number</i>	The size (in bytes) of the MTU value. The configurable MTU range is 1492 through 9216. The default VPLS MTU value is 1500.

Default

dot1q tag - excluded.

ethertype - the configured switch ethertype is used.

number (MTU) - 1500.

Usage Guidelines

This command configures the VPLS parameters. PWs are point-to-point links used to carry VPN traffic between two devices within the VPLS. Each device must be configured such that packets transmitted

between the endpoints are interpreted and forwarded to the local service correctly. The optional `ethertype` keyword may be used to overwrite the Ethertype value for the customer traffic sent across the PW. By default, the configured switch ethertype is used. If configured, the ethertype in the outer 802.1q field of the customer packet is overwritten using the configured ethertype value. The ethertype value is ignored on receipt.

Optionally, the switch can be configured to strip the 802.1q tag before sending packets over the VPLS L2 VPN. This capability may be required to provide interoperability with other vendor products or to emulate port mode services. The default configuration is to include the 802.1q tag.

The `mtu` keyword optionally specifies the MTU value of the VPLS transport payload packet (customer packet). The MTU value is exchanged with VPLS-configured peer nodes. All VPLS peer nodes must be configured with the same MTU value. If the MTU values do not match, PWs cannot be established between VPLS peers. The MTU values are signaled during PW establishment so that endpoints can verify that MTU settings are equivalent before establishing the PW. By default the VPLS MTU is set to 1500. The configurable MTU range is 1492 through 9216. Changing the MTU setting causes established PWs to terminate. VPLS payload packets may be dropped if the VPLS MTU setting is greater than the *MPLS* MTU setting for the PW interface.

**Note**

The maximum MTU value supported depends on the current configuration options. For more information, see [Configuring the Layer 2 VPN MTU](#) in the *Switch Engine 32.3 User Guide*.

Example

The following commands change the various parameters of a particular VPLS:

```
configure vpls vpls1 dot1q ethertype 0x8508
configure vpls vpls1 dot1q ethertype 0x8509 mtu 2500
configure vpls vpls1 dot1q tag exclude mtu 2430
configure vpls vpls1 dot1q mtu 2500
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure vpls add peer

```
configure vpls vpls_name add peer ipaddress {core {full-mesh | primary | secondary} | spoke}
```



Note

This command has been replaced with the following command: . This command is still supported for backward compatibility, but it will be removed from a future release, so we recommend that you start using the new command.

```
configure l2vpn [vpls vpls_name | vpws vpws_name] add peer ipaddress {core {full-mesh | primary | secondary} | spoke}
```

Description

Configures a VPLS or H-VPLS peer for the node you are configuring.

Syntax Description

<i>vpls_name</i>	Specifies the VPLS for which you are configuring a peer.
<i>ipaddress</i>	Specifies the IP address of the peer node.
core	Specifies that the peer is a core node.
full-mesh	Specifies that the peer is a core full-mesh node. This is the default setting if neither the core or spoke options are specified.
primary	Specifies that the peer is an H-VPLS core node and configures a primary H-VPLS connection to that core node.
secondary	Specifies that the peer is an H-VPLS core node and configures a secondary H-VPLS connection to that core node.
spoke	Specifies that the peer is a H-VPLS spoke node.

Default

N/A.

Usage Guidelines

Up to 32 core nodes can be configured for each VPLS. H-VPLS spoke nodes can peer with core nodes. Nodes can belong to multiple VPLS instances. The *ipaddress* parameter identifies the VPLS node that is the endpoint of the VPLS PW.

Core nodes must be configured in a full-mesh with other core nodes. Thus, all core nodes in the VPLS must have a configured PW to every other core node serving this VPLS. By default, the best LSP is chosen for the PW. The underlying LSP used by the PW can be configured by specifying the named LSP using the CLI command `configure l2vpn [vpls vpls_name | vpws vpws_name] peeripaddress [add | delete] mpls lsp_lsp_name .`

Spoke nodes establish up to two point-to-point connections to peer with core nodes. If both primary and secondary peers are defined for a spoke node, the spoke node uses one of the peers for all communications. If both peers are available, the spoke node uses the connection to the primary peer. If the primary peer connection fails, the spoke node uses the secondary peer. If the primary peer later recovers, the spoke node reverts back to using the primary peer.

Example

The following command adds a connection from the local core switch to the core switch at 1.1.1.202:

```
configure vpls vpls1 add peer 1.1.1.202
```

The following command adds a connection from the local core switch to the spoke switch at 1.1.1.201:

```
configure vpls vpls1 add peer 1.1.1.201 spoke
```

The following command adds a primary connection from the local spoke switch to the core switch at 1.1.1.203:

```
configure vpls vpls1 add peer 1.1.1.203 core primary
```

History

This command was first available in ExtremeXOS 11.6.

Support for H-VPLS was first available in ExtremeXOS 12.1.

Platform Availability

This command is available only on the platforms that support *MPLS* as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure vpls delete peer

```
configure vpls vpls_name delete peer [ipaddress | all]
```



Note

This command has been replaced with the command below. This command is still supported for backward compatibility, but it will be removed from a future release, so we recommend that you start using the new command.

```
configure l2vpn [vpls vpls_name | vpws vpws_name] delete peer  
[ipaddress | all]
```

Description

Deletes a VPLS peer from the specified vpls_name.

Syntax Description

<i>vpls_name</i>	Identifies the VPLS within the switch (character string).
<i>ipaddress</i>	Specifies the IP address for the peer node that is the endpoint of the VC-LSP.
all	Deletes all VPLS peers.

Default

N/A.

Usage Guidelines

This command deletes a VPLS peer from the specified *vpls_name*. When the VPLS peer is deleted, VPN connectivity to the VPLS peer is terminated. The **all** keyword may be used to delete all peers associated with the specified VPLS.

Example

The following example removes connectivity to 1.1.1.202 from VPLS1:

```
configure vpls vpls1 delete peer 1.1.1.202
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure vpls delete service

```
configure vpls vpls_name delete service [{vlan} vlan_name | {vman}  
vman_name]
```



Note

This command has been replaced with the following command: `configure l2vpn [vpls vpls_name | vpws vpws_name] delete service [{vlan} vlan_name | {vman} vman_name]`.

This command is still supported for backward compatibility, but it will be removed from a future release, so we recommend that you start using the new command.

Description

Deletes local VPLS service from the specified *vpls_name*.

Syntax Description

<i>vpls_name</i>	Identifies the VPLS interface within the switch (character string).
<i>vlan_name</i>	Logically binds the <i>VLAN</i> to the specified VPLS.
<i>vman_name</i>	Adds the named VMAN to the VPLS.

Default

N/A.

Usage Guidelines

This command deletes the local VPLS service from the specified *vpls_name*. Specifying the *vlan_name* or *vman_name* deletes the service from the VPLS. If there are no services configured for the VPLS, all PWs within the VPLS are terminated from the switch.

Example

The following example removes a service interface from a VPLS:

```
configure vpls vpls1 delete vman vman1
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support *MPLS* as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure vpls health-check vccv

```
configure vpls [vpls_name | all] health-check vccv {interval
interval_seconds} {fault-multiplier fault_multiplier_number}
```



Note

This command has been replaced with the following command: `configure l2vpn [vpls [vpls_name | all] | vpws [vpws_name | all]] health-check vccv {intervalinterval_seconds} {fault-multiplierfault_multiplier_number}`.

This command is still supported for backward compatibility, but it will be removed from a future release, so we recommend that you start using the new command.

Description

Configures the VCCV health check test and fault notification intervals for the specified VPLS instance.

Syntax Description

<i>vpls_name</i>	Identifies the VPLS instance for which health check is to be configured.
all	Specifies that the configuration applies to all VPLS instances on the local node.
<i>interval_seconds</i>	Defines the interval between health check tests. The range is 1 to 10 seconds.
<i>fault_multiplier_number</i>	Specifies how long health check waits before a warning level message is logged. The wait period is the <i>interval_seconds</i> multiplied by the <i>fault_multiplier_number</i> . The <i>fault_multiplier_number</i> range is 2 to 6.

Default

Interval is 5 seconds.

Fault multiplier is 4.

Usage Guidelines

The VCCV health-check configuration parameters can be configured at anytime after the VPLS has been created.

The `show l2vpn {vpls {{vpls_name}} | vpws {{vpws_name}}} {peeripaddress} {detail} | summary` command displays the configured *interval_seconds* and *fault-multiplier_number* values for the VPLS and the VCCV activity state.

Example

The following command configures the health check feature on the VPLS instance myvpls:

```
configure vpls myvpls health-check vccv interval 10 fault-notification 40
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure vpls peer l2pt profile

```
configure {l2vpn} vpls vpls_name peer ipaddress l2pt profile [none |
    profile_name]
```


Description

Configures L2PT profiles on service interfaces.

Syntax Description

l2vpn	Specifies the Layer 2 Virtual Private Network.
vpls <i>vpls_name</i>	Specifies Virtual Private LAN Service over <i>MPLS</i> , and the alphanumeric string identifying the VPLS VPN.
peer <i>ipaddress</i>	Specifies the VPLS peer, and the IPv4 address.
l2pt profile	Specifies Layer 2 protocol tunneling and the L2PT profile for the PW.
none	Specifies that no L2PT profile should be bound to the PW (default).
<i>profile_name</i>	Specifies the L2PT profile to be bound to the PW.

Default

Disabled.

Usage Guidelines

Use this command to configure L2PT profiles on service interfaces.

Example

The following example unbind the L2PT profile from peer 1.1.1.1 of VPLS cust2:

```
configure l2vpn vpls cust2 peer 1.1.1.1 l2pt profile none
```

The following example binds *my_l2pt_prof* with peer 1.1.1.1 of VPLS cust1. *my_l2pt_prof* specifies tunneling actions:

```
configure l2vpn vpls cust1 peer 1.1.1.1 l2pt profile my_l2pt_prof
Error: Tunnel action may be applied only to ports.
```

History

This command was first available in ExtremeXOS 15.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure vpls peer mpls lsp

```
configure vpls vpls_name peer ipaddress [add | delete] mpls lsp lsp_name
```



Note

This command has been replaced with the following command: `configure l2vpn [vpls vpls_name | vpwsvpws vpws_name] peer ipaddress [add | delete] mpls lsp lsp_name .`

This command is still supported for backward compatibility, but it will be removed from a future release, so we recommend that you start using the new command.

Description

Configures a named LSP to be used for the PW to the specified VPLS peer.

Syntax Description

<i>vpls_name</i>	Identifies the VPLS within the switch (character string).
<i>ipaddress</i>	Specifies the IP address for the peer node that is the endpoint of the PW-LSP.
add	Permits addition of up to four RSVP-TE LSPs to the VPLS peer.
delete	Removes the LSP specified by the <i>lsp_name</i> parameter from the PW-LSP aggregation list.
<i>lsp_name</i>	Removes the specified Lsp.

Default

N/A.

Usage Guidelines

This command configures a named LSP to be used for the PW to the specified VPLS peer. The delete keyword removes the LSP specified by the *lsp_name*. If all the named LSPs are deleted to the configured VPLS peer, VPLS attempts to use the best-routed path LSP, if one exists. The delete portion of this command cannot be used to remove a named LSP that was selected by the switch as the best LSP. If no LSPs exist to the VPLS peer, VPN connectivity to the VPLS peer is lost. Currently, the VPLS PW uses only one LSP.

In ExtremeXOS 15.4, this command is modified to display an informational message when multiple transport LSPs are configured for a VPLS PW, when LSP sharing is not enabled. This message is only displayed once per switch boot.

Example

The following examples add and remove a named LSP:

```
configure vpls vpls1 peer 1.1.1.202 add mpls lsp "to-olympic4"
configure vpls vpls1 peer 1.1.1.202 delete mpls lsp "to-olympic4"
configure vpls vpls1 peer 20.20.20.83 add mpls lsp lsp2
```



Note

To share LSPs in HW, use the `enable l2vpn sharing` command.

History

This command was first available in ExtremeXOS 11.6.

This command was modified, in ExtremeXOS 15.4, to display an informational message when multiple transport LSPs are configured for a VPLS PW, and LSP sharing is not enabled.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure vpls peer

```
configure vpls vpls_name peer ipaddress [limit-learning number |
unlimited-learning]
```



Note

This command has been replaced with the following command: `configure l2vpn [vpls vpls_name | vpwsvpws_name] peer ipaddress [limit-learning number | unlimited-learning]`.

This command is still supported for backward compatibility, but it will be removed from a future release, so we recommend that you start using the new command.

Description

Configures the maximum number of MAC SAs (Source Addresses) that can be learned for a given VPLS and peer.

Syntax Description

<i>vpls_name</i>	Identifies the VPLS within the switch (character string).
<i>ipaddress</i>	Specifies the IP address for the peer node that is the endpoint of the PW-LSP.
limit-learning	Specifies a limit to the number of MAC SAs to be learned for the specified VPLS and peer.

<i>number</i>	The maximum number of MAC SAs that can be learned for the specified VPLS and peer.
unlimited-learning	Specifies no limit to the number of MAC SAs to be learned for the specified VPLS and peer.

Default

Unlimited.

Usage Guidelines

This command configures the maximum number of MAC SAs (Source Addresses) that can be learned for a given VPLS and peer. This parameter can only be modified when the specified VPLS is disabled. The **unlimited-learning** keyword can be used to specify that there is no limit. The default value is **unlimited-learning**.

Example

The following example causes no more than 20 MAC addresses to be learned on VPLS1's PW to 1.1.1.202:

```
configure vpls vpls1 peer 1.1.1.202 limit-learning 20
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure vpls snmp-vpn-identifier

```
configure vpls vpls_name snmp-vpn-identifier identifier
```

Description

Configures a [SNMP](#) VPN identifier for traps from the specified VLPLS.

Syntax Description

<i>vpls_name</i>	Specifies the VPLS for which you are configuring the identification string.
<i>identifier</i>	Specifies a text string to identify the VPLS in SNMP traps.

Default

N/A.

Usage Guidelines

None.

Example

The following command configures the identifier vpls1trap for SNMP VPN traps on VPLS vpls1:

```
configure vpls vpls1 snmp-vpn-identifier vpls1trap
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

configure vr add ports

```
configure vr vr-name add ports port_list
```

Description

Assigns a list of ports to the specified VR.

Syntax Description

vr-name	Specifies the name of the VR.
port_list	Specifies the ports to add to the VR.

Default

By default, all ports are assigned to [VR-Default](#).

Usage Guidelines

When a new VR is created, by default, no ports are assigned, no [VLAN](#) interface is created, and no support for any routing protocols is added. Use this command to assign ports to a VR. Since all ports are initially assigned to VR-Default, you might need to delete the desired ports first from the VR where they reside before you add them to the desired VR.

If you plan to assign VR ports to a VLAN, be aware that the ports that you add to a VLAN and the VLAN itself cannot be explicitly assigned to different VRs. When multiple VRs are defined, consider the following guidelines while adding ports to a VR:

- A VLAN can belong (either through explicit or implicit assignment) to only one VR.
- If a VLAN is not explicitly assigned to a VR, then the ports added to the VLAN must be explicitly assigned to a single VR.
- If a VLAN is explicitly assigned to a VR, then the ports added to the VLAN must be explicitly assigned to the same VR or to no VR.
- If a port is added to VLANs that are explicitly assigned to different VRs, the port must be explicitly assigned to no VR.

Example

The following example adds all the ports on slot 2 to the VR "vr-acme":

```
configure vr vr-acme add ports 2:*
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure vr add protocol

```
configure vr vr_name [add | delete] protocol [ospf | ospfv3 | rip |  
ripng | bgp | isis | pim | mpls]
```

Description

Starts a Layer 3 protocol instance for a VR or VRF.

Syntax Description

<i>vr_name</i>	Specifies the name of a VR or a VRF.
protocol	Specifies a Layer 3 protocol that you can add or delete.
name	Specifies the name of a VR or a VRF. The following protocols are supported on VRs: <i>RIP</i> , <i>RIPng</i> , <i>OSPF</i> , <i>OSPFv3</i> , <i>BGP</i> , PIM. IS-IS, and <i>MPLS</i> . The following protocols are supported on VRFs: BGP, OSPFv3.
add	Adds a routing protocol to VRF for PE – CE communication .
delete	Specifies the name of a VR or a VRF.

Default

By default, none of the dynamic protocols are added to a User VR or a VRF.

Usage Guidelines

When a new VR or VRF is created, by default, no ports are assigned, no VLAN interface is created, and no support for any routing protocols is added.

MPLS is the only protocol that you can add to or delete from VR-Default. When MPLS is enabled on a switch, the default configuration adds MPLS to VR-Default. You cannot add or delete any other protocols from VR-Default, and you cannot add or delete any protocols from the other system VRs, VR-Mgmt and VR-Control.



Note

You must delete the MPLS protocol from VR-Default before you can add it to a user VR. MPLS can be active on only one VR within a switch.

When you add a protocol to a VRF, the parent VR starts that protocol, if it was not already running, and adds a protocol instance to support the VRF.



Note

OSPFv3 protocol can be added only to the user VR and non-VPN VRF.

If a previously configured protocol instance is deleted, the CE routes imported from that protocol into the VRF RIB is removed.

Example

The following example starts RIP on the VR "vr-acme":

```
configure vr vr-acme add protocol rip
```

The following example starts a BGP protocol instance for VRF "vr-widget":

```
configure vr vr-widget add protocol bgp
```

History

This command was first available in ExtremeXOS 11.0.

MPLS protocol support was added in ExtremeXOS 12.4.

Support for the OSPFv3 and RIPng protocols on user VRs was added in ExtremeXOS 12.5.

Support for the BGP protocol on VRFs was added in ExtremeXOS 12.6.0-BGP.

Platform Availability

This command is available on ExtremeSwitching 5420, 5520, 5720 series switches.

configure vr delete ports

```
configure vr vr-name delete ports port_list
```

Description

Removes a list of ports from the VR specified.

Syntax Description

<i>vr-name</i>	Specifies the name of the VR.
<i>port_list</i>	Specifies the ports to remove from the VR.

Default

By default, all ports are assigned to VR-Default.

Usage Guidelines

When a new VR is created, by default, no ports are assigned, no VLAN interface is created, and no support for any routing protocols is added. Use this command to remove ports from a VR. Since all ports are initially assigned to VR-Default, you might need to delete the desired ports first from the VR where they reside before you add them to the desired VR.

Example

The following example removes all the ports on slot 2 from the VR "vr-acme":

```
configure vr vr-acme delete ports 2:*
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure vr description

```
configure vr vr_name {description desc_string }
```

Description

Use this command to configure a description for the specified VR or VRF.

Syntax Description

<i>vr_name</i>	Specifies the name of a user VR or a VRF.
<i>desc_string</i>	Specifies a text string to describe the VR. If the text string contains space characters, the entire string must be enclosed with double quotes (" ").

Default

No description.

Usage Guidelines

This command allows you to add comments about a VRF/VR entity. Entering a NULL string on the CLI will unconfigure the description string for the VRF/VR. If the description string has spaces in it, then the string must be enclosed within double quotes (" ").

This text message appears in the `show virtual-router` command display when the command specifies a VR name. For VPN VRFs, this message is returned for a mplsL3VPN MIB query of the MIB variable mplsL3VpnVrfDescription.

Example

The following example configures a description for the VRF "corporate":

```
configure vr corporate description "VRF for the corporate intranet"
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure vrrp group

```
configure vrrp group group_name add [primary-vr | secondary-vr] [{vlan}
    vlan_name vrid vridval | vlan vlan_list {vrid vrid_list}]
configure vrrp group group_name delete [primary-vr | secondary-vr
    [{vlan} vlan_name vrid vridval | vlan vlan_list {vrid vrid_list} |
    all] ]
```

Description

The first version of this command adds a primary VR or secondary VR to a specified group by supplying a VLAN name and VRID. The second version of the command deletes a primary VR or secondary VR from a specified group by supplying a VLAN name and VRID.

Syntax Description

group	Form a group of <i>VRRP</i> VRs to operate in high-scale mode.
<i>group_name</i>	Specifies the VRRP group name.
add	Adds a VR to a VRRP group.
primary-vr	Specifies adding/deleting a primary VR of the VRRP group that sends VRRP advertisement at configured intervals.
secondary-vr	Specifies adding/deleting a secondary VR of the VRRP group that sends VRRP advertisement at a slower rate than the primary VR.
vlan	Specifies a VLAN for the VR.
<i>vlan_name</i>	Specifies the VLAN name for the VR.
vrid	Specifies a VRID for the VR.
<i>vridval</i>	Specifies the VRID for the VR.
delete	Deletes VR(s) from the VRRP group.
all	Specifies that all VRs (secondary and primary) are deleted from the VRRP group.
<i>vlan_list</i>	List of VLAN ID tags (1-4,094).

Default

When adding multiple secondary VRs at once, if no VRIDs are specified, all VRs configured on the specified VLANs are added to the group.

Example

The following example adds a primary VR VLAN "v1", VRID "1" for VRRP group "ExtremeNet":

```
configure vrrp group ExtremeNet add primary-vr vlan v1 vrid 1
```

The following example adds a set of VRRP VRs configured on VLANs having VLAN IDs ranging from 11 to 20. Out of all of the VRs configured on these VLANs only VRs with VRID ranging from 1 to 2 are added to the VRRP group:

```
Configure vrrp group ExtremeNet add secondary-vr vlan 11-20 vrid 1-2
```

The following example adds all VRs configured on given VLANs to the group as secondary VRs:

```
configure vrrp group ExtremeNet add secondary-vr vlan 11-20
```

History

This command was first available in ExtremeXOS 22.2.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the VRRP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure vrrp fabric-routing

```
configure vrrp vlan [vlan_name | vlan_list] vrid [vridval | vrid_list]  
    {group group_name }fabric-routing [on | off]
```

Description

This command configures fabric routing.

Syntax Description

group	Specifies VRRP VRs information that form the group.
<i>group_name</i>	Name of the specific VRRP group.
fabric-routing	Configures fabric routing on all members of the group.
on	Enables fabric routing capability.
off	Disables fabric routing capability.
<i>port_list</i>	Port list separated by a comma or -.
<i>vlan_name</i>	Specifies the name of a VRRP VLAN .
<i>vlan_list</i>	VLAN list (1-4,094).
<i>vridval</i>	Specifies the VRID for the VRRP instance. To display the configured VRRP router instances, enter the show vrrp command.
<i>vrid_list</i>	List of virtual router IDs (1-255).

Default

N/A

Usage Guidelines

This configuration can be present on all VRRP routers, regardless of the VRRP state of the router. Fabric routing is enabled only when the VRRP router is in backup state.

You need to configure fabric routing on all members of group when a member's VRID is reused in another group.

Example

The following command turns on fabric routing capability on all VR members of the group "ExtremeNet":

```
configure vrrp group ExtremeNet fabric-routing on
```

History

This command was first available in ExtremeXOS 22.2.

VLAN and VR list options added in ExtremeXOS 22.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the VRRP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure vrrp vlan vrid accept-mode

```
configure vrrp vlan vlan_name vrid vridval accept-mode [on | off]
```

Description

Configures a backup [VRRP](#) router instance to accept or reject packets addressed to the IP address owner when operating as the VRRP master.

Additionally, this command provides capability for switches to configure the VRRP virtual IP as NTP server address.

Syntax Description

<i>vlan_name</i>	Specifies the name of a VRRP VLAN .
<i>vridval</i>	Specifies the VRID of a VRRP instance. To display the configured VRRP router instances, enter the show vrrp command.
on	Specifies that the VRRP instance is to accept packets addressed to the IP address owner.
off	Specifies that the VRRP instance not accept packets addressed to the IP address owner. Note: Ping packets are accepted, regardless of the configuration for this command.

Default

Off.

Usage Guidelines

When a backup VRRP router operates as master, it accepts VRRP traffic and routes traffic. The backup router in master mode also accepts ping packets and IPv6 neighbor solicitations and advertisements. However, because the backup router is not the IP address owner, the default configuration rejects all other traffic addressed to the IP address owner.

If your network requires that a backup VRRP router in master mode accept all traffic addressed to the IP address owner, use this command to configure accept-mode on.

In the ExtremeXOS 15.3 release, NTP VRRP Virtual IP support is added. This feature allows you to configure the VRRP virtual IP as NTP server address. The NTP server when configured on the VRRP master will listen on the physical and virtual IP address for NTP clients. For this feature to work correctly, you need to enable accept mode in VRRP. Enabling accept mode allows the switch to process non-ping packets that have a destination IP set to the virtual IP address.

Example

The following example configures a backup VRRP router in master mode to accept packets addressed to the IP address owner:

```
configure vrrp vlan vlan-1 vrid 1 accept-mode on
```

History

This command was first available in ExtremeXOS 12.7.

NTP VRRP Virtual IP support was added in ExtremXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the VRRP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure vrrp vlan vrid add ipaddress

```
configure vrrp vlan [vlan_name | vlan_id] vrid vridval add ipaddress
```

Description

Associates a virtual IP address with a specific *VRRP* instance.

Syntax Description

<i>vlan_name</i>	Specifies the name of a VRRP <i>VLAN</i> .
<i>vlan_id</i>	VLAN ID tag (1-4,094).
<i>vridval</i>	Specifies the VRID of a VRRP instance. To display the configured VRRP router instances, enter the show vrrp command.
<i>ipaddress</i>	Specifies a virtual IPv4 or IPv6 address to be assigned to the VRRP instance.

Default

N/A.

Usage Guidelines

Each VRRP instance is identified by an ID number, VLAN name, and virtual IP address. When two or more routers are configured with the same VRRP ID number, VLAN name, and virtual IP address, the routers with matching parameters are all part of the same VRRP instance. One router within the instance will become the VRRP instance master, and the others will become backup routers for the VRRP instance.

Most routers within a VRRP instance will have a virtual IP address that is different from the actual IP addresses configured on the router. If the virtual IP address for a VRRP instance matches an IP address configured on a host router, the VRRP instance is known as the IP address owner. On the IP address owner, the VRRP instance priority defaults to 255, and by default, the IP address owner becomes the VRRP master when VRRP is enabled.



Note

There is no requirement to configure an IP address owner within a VRRP instance.

Before each VRRP router is enabled, it must be configured with at least one virtual IPv4 or IPv6 address. You can repeat this command to add additional virtual IP addresses to the VRRP router. If a virtual IPv4 address is added to a VRRP router, you cannot later add a virtual IPv6 address. Similarly, if a virtual IPv6 address is added to a VRRP router, you cannot later add a virtual IPv4 address.

Each IPv6 VRRP instance is associated with one and only one virtual link local address, which serves as the source IP address for subsequent router announcement packets generated by the master VRRP router. The virtual link local address can be explicitly configured or generated automatically. One way to explicitly configure the virtual link local address is to add it to the virtual IP address list with this command.

Example

The following example associates virtual IPv4 address 10.1.2.3 to VRRP router instance 1:

```
configure vrrp vlan vlan-1 vrid 1 add 10.1.2.3
```

The following example associates virtual IPv6 address 2001:db8::3452 to VRRP router instance 2:

```
configure vrrp vlan vlan-1 vrid 2 add 2001:db8::3452
```

History

This command was first available in ExtremeXOS 10.1.

Support for IPv6 addresses was added in ExtremeXOS 12.7.

The `vlan_id` option was added in ExtremeXOS 22.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the VRRP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure vrrp vlan vrid add track-iproute

```
configure vrrp vlan vlan_name vrid vridval add track-iproute ipaddress/  
masklength
```

Description

Creates a tracking entry for the specified route. When this route becomes unreachable, this entry is considered to be failing.

Syntax Description

<i>vlan_name</i>	Specifies the name of a <i>VRRP VLAN</i> .
<i>vridval</i>	Specifies the VRID of the target VRRP instance. To display the configured VRRP router instances, enter the show vrrp command.
<i>ipaddress</i>	Specifies the IPv4 or IPv6 prefix of the route to track.
<i>masklength</i>	Specifies the length of the route's prefix.

Default

N/A.

Usage Guidelines

The route specified in this command might not exist in the IP routing table. When you create the entry for a route, an immediate VRRP failover might occur.



Note

VRRP tracking is not supported on *MPLS* LSPs.

Example

The following command enables IP route failure tracking for routes to the specified subnet:

```
configure vrrp vlan vlan-1 vrid 1 add track-iproute 3.1.0.0/24
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the VRRP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure vrrp vlan vrid add track-ping

```
configure vrrp vlan vlan_name vrid vridval add track-ping ipaddress  
      frequency seconds miss misses [success successes]
```

Description

Creates a tracking entry for the specified IP address. The entry is tracked using pings to the IP address, sent at the specified frequency.

Syntax Description

<i>vlan_name</i>	Specifies the name of a <i>VRRP VLAN</i> .
vrid <i>vridval</i>	Specifies the VRID of the target VRRP instance. To display the configured VRRP router instances, enter the show vrrp command.
<i>ipaddress</i>	Specifies the IPv4 or IPv6 address to be tracked.
frequency <i>seconds</i>	Specifies the number of seconds between pings to the target IP address. The range is 1 to 600 seconds.
miss <i>misses</i>	Specifies the number of misses allowed before this entry is considered to be failing. The range is 1 to 255 pings.
success <i>successes</i>	Sets how many ping successes are required for tracking success. Range is 1–255. (Default is 10 × <i>misses</i> .)

Default

If the number of successes is not specified, the default is ten times the number of misses specified.

Usage Guidelines

Adding an entry with the same IP address as an existing entry causes the new values to overwrite the existing entry's frequency and miss number.

Example

The following command enables ping tracking for the external gateway at 3.1.0.1, pinging every 3 seconds, and considering the gateway to be unreachable if no response is received to 5 consecutive pings:

```
configure vrrp vlan vlan-1 vrid 1 add track-ping 3.1.0.1 frequency 3 miss 5
```

History

This command was first available in ExtremeXOS 10.1.

The **success** option was added in ExtremeXOS 22.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the VRRP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document./ph ">

configure vrrp vlan vrid add track-vlan

```
configure vrrp vlan vlan_name vrid vridval add track-vlan
target_vlan_name
```

Description

Configures a VRRP VLAN to track port connectivity to a specified VLAN. When this VLAN is in the down state, this entry is considered to be failing.

Syntax Description

<i>vlan_name</i>	Specifies the name of a VRRP VLAN.
<i>vridval</i>	Specifies the VRID of the target VRRP instance. To display the configured VRRP router instances, enter the show vrrp command.
<i>target_vlan_name</i>	Specifies the name of the VLAN to track.

Default

N/A.

Usage Guidelines

Up to eight VLANs can be tracked.

Deleting a tracked VLAN does not constitute a failover event for the VRRP VLAN tracking it, and the tracking entry is deleted.

Example

The following command enables VRRP VLAN vlan-1 to track port connectivity to VLAN vlan-2:

```
configure vrrp vlan vlan-1 vrid 1 add track-vlan vlan-2
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the VRRP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure vrrp vlan vrid add virtual-link-local

```
configure vrrp vlan vlan_name vrid vridval add virtual-link-local
vll_addr
```

Description

Specifies a virtual IPv6 link local address for the *VRRP* router instance.

Syntax Description

<i>vlan_name</i>	Specifies the name of a VRRP <i>VLAN</i> .
<i>vridval</i>	Specifies the VRID of a VRRP instance. To display the configured VRRP router instances, enter the show vrrp command.
<i>vll_addr</i>	Specifies a virtual link local address to be assigned to the VRRP instance.

Usage Guidelines

Each IPv6 VRRP instance is associated with one and only one virtual link local address, which serves as the source IP address for subsequent router announcement packets generated by the master VRRP router. The virtual link local address can be explicitly configured or generated automatically.

One way to explicitly configure the virtual link local address is to add it to the virtual IP address list with this command. The new link local address must match the FE80::/64 subnet, and it must match the address in use on all other router in this VRRP instance.

If no virtual link local address is configured, an appropriate address is generated automatically.



Note

If an IPv4 address has been added to a VRRP router, you cannot later add any IPv6 address, so you cannot add a link local address.

Example

The following example associates virtual IPv6 link local address fe80::1111 to VLAN vlan-1:

```
configure vrrp vlan vlan-1 vrid 1 add virtual-link-local fe80::1111
```

History

This command was first available in ExtremeXOS 12.7.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the VRRP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure vrrp vlan vlan_name vrid vridval advertisement-interval

```
configure vrrp vlan vlan_name vrid vridval advertisement-interval
interval [{seconds} | centiseconds]
```

Description

Configures the time between [VRRP](#) advertisements in seconds or centiseconds.

Syntax Description

<i>vlan_name</i>	Specifies the name of a VRRP VLAN .
<i>vridval</i>	Specifies the VRID of a VRRP instance. To display the configured VRRP router instances, enter the show vrrp command.
<i>interval</i>	Specifies an interval value for the time between advertisements. The range is 1 through 40 seconds or 10 through 4095 centiseconds.
seconds	Specifies that the interval value is in seconds. If you do not specify seconds or centiseconds, the interval value is applied as seconds.
centiseconds	Specifies that the interval value is in centiseconds.

Default

The advertisement interval is 1 second.

Usage Guidelines

The advertisement interval specifies the interval between advertisements sent by the master router to inform the backup routers that its alive. You must use whole integers when configuring the advertisement interval.

An extremely busy CPU can create a short dual master situation. To avoid this, increase the advertisement interval.



Note

The milliseconds keyword is replaced by the centiseconds keyword, but the milliseconds keyword is still recognized to support existing configurations and scripts. Any values specified in milliseconds are converted to centiseconds. All new configurations and scripts should specify the interval in either seconds or centiseconds. The maximum value for an interval specified in seconds is 40. However, the software supports older configurations and scripts that specify values up to 255, which were supported prior to ExtremeXOS Release 12.7.

To view your VRRP configuration, including the configured advertisement interval, use one of the following commands:

- `show vrrp {virtual-router {vr-name}} {detail}`
- `show vrrp vlan vlan_name {stats}`

If you enter a number that is out of the seconds or centiseconds range, the switch displays an error message. For example, if the interval value is set to 999 and the centiseconds keyword is missing, the switch displays an error message similar to the following:

```
configure vrrp blue vrid 250 advertisement-interval 999 Error:
Advertisement interval must be between 1 and 255 seconds. 999 out of
range
```

Example

The following command configures the advertisement interval for 15 seconds:

```
configure vrrp vlan vrrp-1 vrid 1 advertisement-interval 15
```

The following command configures the advertisement interval for 200 centiseconds:

```
configure vrrp vlan vrrp-1 vrid 1 advertisement-interval 200 centiseconds
```

History

This command was first available in ExtremeXOS 10.1.

The milliseconds and seconds keywords were added in ExtremeXOS 11.5.

The centiseconds keyword replaced the milliseconds keyword, and the maximum value for intervals specified in seconds was reduced to 40 in ExtremeXOS 12.7.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the VRRP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure vrrp vlan vrid delete track-iproute

```
configure vrrp vlan vlan_name vrid vridval delete track-iproute
ipaddress/masklength
```

Description

Deletes a tracking entry for the specified route.

Syntax Description

<i>vlan_name</i>	Specifies the name of a VRRP VLAN .
<i>vridval</i>	Specifies the VRID of the target VRRP instance. To display the configured VRRP router instances, enter the show vrrp command.
<i>ipaddress</i>	Specifies the IPv4 or IPv6 prefix of the route.
<i>masklength</i>	Specifies the length of the route's prefix.

Default

N/A.

Usage Guidelines

Deleting a tracking entry while VRRP is enabled causes the VRRP VRs state to be re-evaluated for failover.

Example

The following command disables tracking of routes to the specified subnet for VLAN vlan-1:

```
configure vrrp vlan vlan-1 vrid 1 delete track-iproute 3.1.0.0/24
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the VRRP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure vrrp vlan vrid delete track-ping

```
configure vrrp vlan vlan_name vrid vridval delete track-ping ipaddress
```

Description

Deletes a tracking entry for the specified IP address.

Syntax Description

<i>vlan_name</i>	Specifies the name of a VRRP VLAN .
<i>vridval</i>	Specifies the VRID of the target VRRP instance. To display the configured VRRP router instances, enter the show vrrp command.
<i>ipaddress</i>	Specifies the IP address to be tracked.

Default

N/A.

Usage Guidelines

Deleting a tracking entry while VRRP is enabled causes the VRRP VRs state to be re-evaluated for failover.

A VRRP node with a priority of 255 might not recover from a ping-tracking failure if there is a Layer 2 switch between it and another VRRP node. In cases where a Layer 2 switch is used to connect VRRP nodes, we recommend that those nodes have priorities of less than 255.

Example

The following command disables ping tracking for the external gateway at 3.1.0.1:

```
configure vrrp vlan vlan-1 vrid 1 delete track-ping 3.1.0.1
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the VRRP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure vrrp vlan vrid delete track-vlan

```
configure vrrp vlan vlan_name vrid vridval delete track-vlan  
                        target_vlan_name
```

Description

Deletes the tracking of port connectivity to a specified [VLAN](#).

Syntax Description

<i>vlan_name</i>	Specifies the name of a <i>VRRP</i> VLAN.
<i>vridval</i>	Specifies the VRID of the VRRP instance. To display the configured VRRP router instances, enter the <code>show vrrp</code> command.
<i>target_vlan_name</i>	Specifies the name of the tracked VLAN.

Default

N/A.

Usage Guidelines

Deleting a tracking entry while VRRP is enabled causes the VRRP VRs state to be re-evaluated for failover.

Example

The following command disables the tracking of port connectivity to VLAN vlan-2:

```
configure vrrp vlan vlan-1 vrid 1 delete track-vlan vlan-2
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the VRRP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure vrrp vlan vrid delete ipaddress

```
configure vrrp vlan [vlan_name | vlan_id] vrid vridval delete ipaddress
```

Description

Deletes a virtual IPv4 or IPv6 address from a specific *VRRP* router.

Syntax Description

<i>vlan_name</i>	Specifies the name of a VRRP <i>VLAN</i> .
<i>vlan_id</i>	VLAN ID tag (1-4,094).

<i>vridval</i>	Specifies the VRID of the VRRP instance. To display the configured VRRP router instances, enter the show vrrp command.
<i>ipaddress</i>	Specifies the virtual IP address to be deleted from the VRRP instance. This is common for IPv4/IPv6 addresses.

Usage Guidelines

When a VRRP router is enabled, it must have at least one virtual IP address. When the VRRP router is not enabled, there are no restrictions on deleting the IP address.

Example

The following command removes IP address 10.1.2.3 from VLAN vlan-1:

```
configure vrrp vlan vlan-1 vrid 1 delete 10.1.2.3
```

History

This command was first available in ExtremeXOS 10.1.

The *vlan_id* option was added in ExtremeXOS 22.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the VRRP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure vrrp vlan vrid dont-preempt

```
configure vrrp vlan vlan_name vrid vridval dont-preempt
```

Description

Specifies that a higher priority backup router does not preempt a lower priority master.

Syntax Description

<i>vlan_name</i>	Specifies the name of a <i>VRRP VLAN</i> .
<i>vridval</i>	Specifies the VRID of a VRRP instance. To display the configured VRRP router instances, enter the show vrrp command.

Default

The default setting is preempt.

Usage Guidelines

The preempt mode controls whether a higher priority backup router preempts a lower priority master. dont-preempt prohibits preemption. The router that owns the virtual IP address always preempts, independent of the setting of this parameter.

Example

The following command disallows preemption:

```
configure vrrp vlan vlan-1 vrid 1 dont-preempt
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the VRRP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure vrrp vlan vrid host-mobility

```
configure vrrp [{vlan} [vlan_name | vlan_list] vrid [vridval | vrid_list
| all]] host-mobility [{on | off} {exclude-ports [add | delete]
port_list}]
```

Description

All instances of VRRP have host-mobility off by default. Configuring host-mobility to “on” state starts ARP route learning. By default, all ports perform the route learning. Configuring host-mobility excluded-ports will disable the route learning on the port list provided. All ports of the VRRP VLAN that are connected to another router should be excluded. If ports are not excluded, routes are created for devices as if they are directly connected and this may cause traffic to take a longer route.

Syntax Description

host-mobility	Exportable Host Route learning via ARP/ND on the specified VLAN and VRID.
on	Advertise host routes for hosts learned via ARP/ND.
off	Do not advertise host routes for hosts learned via ARP/ND.
exclude-ports	Exclude ports from host-mobility route learning (Default: no ports are excluded).
add	Add ports to host-mobility exclude list; host-mobility routes will not be learned on the ports.
delete	Delete ports from host-mobility exclude list.

<code>port_list</code>	Port list separated by a comma or -“.
<code>vlan_name</code>	Specifies the name of a VRRP VLAN .
<code>vlan_list</code>	VLAN list (1-4,094).
<code>vridval</code>	Specifies the VRID for the VRRP instance. To display the configured VRRP router instances, enter the show vrrp command.
<code>vrid_list</code>	List of virtual router IDs (1-255).
all	Selects all VRRP virtual routers.

Default

Off.

Usage Guidelines

Configuring host-mobility excluded-ports will disable the route learning on the port list provided. All ports that are connected to another router should be excluded. If ports are not excluded, routes will be created for devices as if they are directly connected and may cause traffic to take a longer route.

Example

```
configure vrrp vlan vlan1 vrid 1 host-mobility on excluded-ports add 1,10
```

History

This command was first available in ExtremeXOS 21.1.

VLAN and VR list options added in ExtremeXOS 22.3.

The **all** option was added in ExtremeXOS 22.4.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the VRRP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure vrrp vlan vrid ipv4 checksum

```
configure vrrp {vlan} vlan_name vrid vridval ipv4 checksum [include-pseudo-header | exclude-pseudo-header]
```

Description

This command allows you to eliminate the pseudo header for VRRPv3 IPv4 Checksum calculation.

Default

Include.

Example

```
configure vrrp vlan "v1" vrid 1 ipv4 checksum exclude-pseudo-header
```

History

This command was first available in ExtremeXOS 16.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the VRRP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure vrrp vlan vrid preempt

```
configure vrrp vlan vlan_name vrid vridval preempt {delay seconds}
```

Description

Specifies that a higher priority backup router preempts a lower priority master.

Syntax Description

<i>vlan_name</i>	Specifies the name of a <u>VRRP VLAN</u> .
<i>vridval</i>	Specifies the VRID for a VRRP instance. To display the configured VRRP router instances, enter the show vrrp command.
<i>seconds</i>	Specifies a preempt delay period in seconds. The value range is 1 to 3600 seconds, or 0, which selects the original preempt delay period.

Default

Preempt enabled.

Delay configuration: 0.

Usage Guidelines

The preempt option enables a higher-priority backup router to preempt a master with a lower priority. When a VRRP enabled router receives a lower priority VRRP advertisement and preemption is enabled,

the higher-priority VRRP enabled router takes over as master. The new master starts sending VRRP advertisements and the old, lower-priority master relinquishes mastership.

**Note**

The router that owns the virtual IP address always preempts, independent of the setting of this parameter.

When a VRRP enabled router preempts the master, it does so in one of the following ways:

- If the preempt delay timer is configured for between 1 and 3600 seconds and the lower-priority master is still operating, the router preempts the master when the timer expires.
- If the preempt delay timer is configured for 0, the router preempts the master after 3 times the hello interval.
- If the higher priority router stops receiving advertisements from the current master for 3 times the hello interval, it takes over mastership immediately.

**Note**

The preempt feature can be disabled with the `configure vrrp vlan vrid dont-preempt` command.

Example

The following command allows preemption:

```
configure vrrp vlan vlan-1 vrid 1 preempt
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the VRRP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure vrrp vlan vrid priority

```
configure vrrp vlan [vlan_name | vlan_list] vrid [vridval | vrid_list]  
           priority priorityval
```

Description

Configures the priority value of a VRRP router instance.

Syntax Description

<i>vlan_name</i>	Specifies the name of a VRRP VLAN .
<i>vlan_list</i>	VLAN list (1-4,094).
<i>vridval</i>	Specifies the VRID for the VRRP instance. To display the configured VRRP router instances, enter the show vrrp command.
<i>vrid_list</i>	List of virtual router IDs (1-255).
<i>priorityval</i>	Specifies the priority value of the router. The default is 100. The priority range is 1-255.

Default

The default priority is 100.

Usage Guidelines

This command changes the priority of a VRRP router. If the VRRP router is the IP address owner (which means that the VRRP router IP address matches the VRRP VLAN IP address), the priority is 255 and cannot be changed. If the VRRP router is not the IP address owner, the priority can be changed to values in the range of 1 to 254.

To change the priority of the IP address owner or to make a different VRRP router the IP address owner, disable VRRP and reconfigure the affected switches to use VRRP router addresses that support the priorities you want to assign.

Example

The following command configures a priority of 150 for VLAN vrrp-1:

```
configure vrrp vlan vrrp-1 vrid 1 priority 150
```

History

This command was first available in ExtremeXOS 10.1.

VLAN and VR list options added in ExtremeXOS 22.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the VRRP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure vrrp vlan vrid track-mode

```
configure vrrp vlan vlan_name vrid vridval track-mode [all | any]
```

Description

Defines the conditions under which the router automatically relinquishes master status when the tracked entities fail.

Syntax Description

<i>vlan_name</i>	Specifies the name of a <i>VRRP VLAN</i> .
<i>vridval</i>	Specifies the VRID for the VRRP instance. To display the configured VRRP router instances, enter the show vrrp command.
all	Specifies that the mastership is relinquished when one of the following events occur: All of the tracked VLANs fail All of the tracked routes fail All of the tracked PINGs fail
any	Specifies that the mastership is relinquished when any of the tracked VLANs, routes, or PINGs fail.

Default

The default setting is all.

Usage Guidelines

None.

Example

The following command configures the track mode to any:

```
configure vrrp vlan vrrp-1 vrid 1 track-mode any
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the VRRP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure vrrp vlan vrid version

```
configure vrrp vlan vlan_name vrid vridval version [v3-v2 | v3 | v2]
```

Description

Selects the *VRRP* version to apply to the VRRP router instance.

Syntax Description

<i>vlan_name</i>	Specifies the name of a VRRP <i>VLAN</i> .
<i>vridval</i>	Specifies the VRID for the VRRP instance. To display the configured VRRP router instances, enter the <code>show vrrp</code> command.
v3-v2	Specifies VRRP v3 with VRRP v2 compatibility.
v3	Selects VRRP v3.
v2	Specifies VRRP v2.

Default

VRRP v3 with VRRP v2 compatibility.



Note

Configurations created by earlier ExtremeXOS software releases have an implied version of v2. If the configuration is subsequently saved, the version is explicitly set to v2.

Usage Guidelines

None.

Example

The following command configures the VRRP router instance to use VRRP v3 only:

```
configure vrrp vlan vrrp-1 vrid 1 version v3
```

History

This command was first available in ExtremeXOS 12.7.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the VRRP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

configure web http access-profile

```
configure web http access-profile [[[add rule ] [first | [[before |
after] previous_rule]]] | delete rule | none ]
```

Description

Configures HTTP to use an *ACL* rule for access control.

Syntax Description

add	Specifies that an ACL rule is to be added to the website.
<i>rule</i>	Specifies an ACL rule.
first	Specifies that the new rule is to be added before all other rules.
before	Specifies that the new rule is to be added before a previous rule.
after	Specifies that the new rule is to be added after a previous rule.
<i>previous_rule</i>	Specifies an existing rule in the application.
delete	Specifies that one particular rule is to be deleted.
none	Specifies that all the rules or a policy file is to be deleted.

Default

N/A.

Usage Guidelines

You must be logged in as administrator to configure HTTP parameters.

Use this command to restrict HTTP access by adding an ACL rule to the HTTP application. Once an ACL is associated with HTTP, all the packets that reach a HTTP module are evaluated with this ACL and appropriate action (permit or deny) is taken, as is done using policy files.

The permit or deny counters are also updated accordingly regardless of whether the ACL is configured to add counters. To display counter statistics, use the [tftp put](#) on page 3377 http command.

Only the following match conditions and actions are copied to the client memory. Others that may be in the rule are not copied.

Match conditions

- Source-address—IPv4 and IPv6

Actions

- Permit
- Deny

When adding a new rule, use the first, before, and after previous_rule parameters to position it within the existing rules.

If the *SNMP* traffic does not match any of the rules, the default behavior is permit. To deny SNMP traffic that does not match any of the rules, add a deny all rule at the end of the rule list.

Example

The following example copies the ACL rule, DenyAccess to the HTTP application in first place:

```
configure web http access-profile add DenyAccess first
```


The following example removes the association of the ACL rule DenyAccess from the HTTP application:

```
configure web http access-profile delete DenyAccess
```

The following example removes the association of all ACL rules from the HTTP application:

```
configure web http access-profile none
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure xml-notification target add/delete

```
configure xml-notification target target [add | delete] module
```

Description

Adds or deletes an ExtremeXOS module to or from the Web server target.

Syntax Description

<i>target</i>	Specifies an alpha numeric string that identifies the configured target.
<i>module</i>	Specifies the name of the ExtremeXOS module.

Default

N/A.

Usage Guidelines

Use the add option to attach a module to the Web server target in order to receive events from that application and send them to the targeted Web server. There is no limitation to the number of modules that can be attached.

Only Identity Management and EMS are supported targets.

Use the delete option to detach ExtremeXOS modules from the Web server target in order to stop receiving events from that module.

Example

The following command deleted the target test2 from EMS:

```
configure xml-notification target test2 ems
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure xml-notification target

```
configure xml-notification target target [url url {vr vr_name} | user
    [none | user] | [encrypted-auth encrypted-auth] | [queue-size queue-
    size]]
```

Description

Configures the Web server target in the XML client.

Syntax Description

<i>target</i>	Specifies an alpha numeric string that identifies the configured target.
<i>url</i>	Specifies the Web server URL.
<i>vr_name</i>	Specifies the virtual router over which the XML client process can connect to a Web server to send push notifications.
<i>user</i>	Specifies the alpha numeric string identifying the Web server user.
<i>encrypted-auth</i>	Specifies the encrypted user authentication string.
<i>queue-size</i>	Specifies in numeric format, the size of the buffer that stores incoming events from ExtremeXOS software.

Default

N/A.

Usage Guidelines

Use this command to configure the Web server target in XML client process.

Example

The following command configures the target target2 for the user admin:

```
configure xml-notification target target2 user admin
```

History

This command was first available in ExtremeXOS 12.4.

The virtual router option was added in ExtremeXOS 12.4.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

configure l2pt encapsulation dest-mac

```
configure l2pt encapsulation dest-mac mac_address
```

Description

Configures the destination address MAC that L2PT encapsulated packets use.

Syntax Description

encapsulation	Specifies Layer 2 protocol tunneling encapsulation.
dest-mac	Specifies the destination MAC address to use for encapsulated PDUs.
<i>mac_addr</i>	Specifies the MAC address.

Default

Usage Guidelines

NA

Example

The following example sets the L2PT destination address MAC to 01:00:00:01:01:02:

```
configure l2pt encapsulation dest-mac 01:00:00:01:01:02
```

History

This command was first available in ExtremeXOS 15.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

cp

```
cp old_name new_name
```

Description

Copies a file from the specified file system or relative to the current working directory to another file on the specified file system or relative to the current working directory.

Syntax Description

<code>old_name</code>	Specifies the name of the file that you want to copy.
<code>new_name</code>	Specifies the name of the newly copied configuration or policy file.

Default

N/A.

Usage Guidelines

Use this command to copy a file from the specified file system, or relative to the current working directory to another file on the specified file system, or relative to the current working directory. If you provide a different name, the new file can be created in the same directory as the existing file.

When you copy a configuration or policy file, remember the following:

- XML-formatted configuration files have a `.cfg` file extension. The switch only runs `.cfg` files.
- ASCII-formatted configuration files have a `.xsf` file extension. For more information, see *Software Upgrade and Boot Options* in the [Switch Engine 32.3 User Guide](#).
- Policy files have a `.pol` file extension.
- Core dump files have a `.gz` file extension.

When you copy a configuration or policy file from the system, make sure you specify the appropriate file extension. For example, when you want to copy a policy file, specify the file name and `.pol`.

When you copy a file on the switch, the switch displays a message similar to the following:

```
Copy config test.cfg to config test1.cfg on switch? (y/n)
```

Type `y` to copy the file. Type `n` to cancel this process and not copy the file.

When you type `y`, the switch copies the file with the new name and keeps a backup of the original file with the original name. After the switch copies the file, use the `ls` command to display a complete list of files. In this example, the switch displays the original file named `test.cfg` and the copied file named `test_rev2.cfg`.

The following is sample output from the `ls` command:

```
...
-rw-r--r--  1 root    root      100980 Sep 23 09:16 test.cfg
-rw-r--r--  1 root    root      100980 Oct 13 08:47 test_rev2.cfg
...
```

When you enter `n`, the switch displays a message similar to the following:

```
Copy cancelled.
```

Case-sensitive Filenames

File names are case-sensitive. In this example, you have a configuration file named `Test.cfg`. If you attempt to copy the file with the incorrect case, for example `test.cfg`, the switch displays a message similar to the following:

```
Error: cp: /config/test.cfg: No such file or directory
```

Since the switch is unable to locate `test.cfg`, the file is not copied.

Local File Name Character Restrictions

When specifying a local file name, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z).
- Numerals (0-9).
- Period (.).
- Dash (-).
- Underscore (_).

Internal Memory and Core Dump Files

Core dump files have a `.gz` file extension. The file name format is: `core.process-name.pid.gz` where `process-name` indicates the name of the process that failed and `pid` is the numerical identifier of that process.

By making a copy of a core dump file you can easily compare new debug information with the old file if needed.

If you have a USB 2.0 storage device installed, you can copy the core dump file to that location. To copy files to a USB device, specify the file path `/usr/local/ext`.

For information about configuring and sending core dump information to the internal memory card, see the `configure debug core-dumps [off | directory_path]` and `save debug tracefiles directory_path` commands.

For more detailed information about core dump files, see *Troubleshooting* in the [Switch Engine 32.3 User Guide](#).

Example

The following example makes a copy of a configuration file named `test.cfg` and gives the copied file a new name of `test_rev2.cfg`:

```
# cp test.cfg test_rev2.cfg
```

The following example makes a copy of a configuration file named `primary.cfg` on the switch and stores the copy on the removable storage device with the same name, `primary.cfg`:

```
# cp primary.cfg /usr/local/ext
```

The above command performs the same action as entering:

```
# cp primary.cfg /usr/local/ext
```

Or

```
# cp primary.cfg /usr/local/ext/primary.cfg
```

History

This command was first available in ExtremeXOS 11.0.

The **memorycard** option was added in ExtremeXOS 11.1.

The **internal-memory** option was added in ExtremeXOS 11.4.

Support for USB 2.0 storage devices was added in ExtremeXOS 12.5.3.

Pathname support was added in ExtremeXOS 15.5.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

create access-list

```
create access-list dynamic_rule conditions actions {non_permanent}
```

Description

Creates a dynamic *ACL*.

Syntax Description

<i>dynamic_rule</i>	Specifies the dynamic ACL name. The name can be from 1-32 characters long.
<i>conditions</i>	Specifies the match conditions for the dynamic ACL.
<i>actions</i>	Specifies the actions for the dynamic ACLs.
non_permanent	Specifies that the ACL is not to be saved.

Default

By default, ACLs are permanent.

Usage Guidelines

This command creates a dynamic ACL rule. Use the `configure access-list add` command to apply the ACL to an interface.

The conditions parameter is a quoted string of match conditions, and the actions parameter is a quoted string of actions. Multiple match conditions or actions are separated by semi-colons. A complete listing of the match conditions and actions is in the *ACLs* section of the [Switch Engine 32.3 User Guide](#).

Dynamic ACL rule names must be unique, but can be the same as used in a policy-file based ACL. Any dynamic rule counter names must be unique. For name creation guidelines and a list of reserved names, see *Object Names* in the [Switch Engine 32.3 User Guide](#).

By default, ACL rules are saved when the save command is executed, and persist across system reboots. Configuring the optional keyword non-permanent means the ACL will not be saved.

Example

The following example creates a dynamic ACL that drops all *ICMP* echo-request packets on the interface:

```
create access-list icmp-echo "protocol icmp;icmp-type echo-request" "deny"
```

The created dynamic ACL will take effect after it has been configured on the interface. The previous example creates a dynamic ACL named icmp-echo that is equivalent to the following ACL policy file entry:

```
entry icmp-echo {
  if {
    protocol icmp;
    icmp-type echo-request;
  } then {
    deny;
  }
}
```

The following example creates a dynamic ACL that accepts all the UDP packets from the 10.203.134.0/24 subnet that are destined for the host 140.158.18.16, with source port 190 and a destination port in the range of 1200 to 1250:

```
create access-list udpacl "source-address 10.203.134.0/24;destination-address
140.158.18.16/32;protocol udp;source-port 190;destination-port 1200 - 1250;" "permit"
```

The previous example creates a dynamic ACL entry named udpacl that is equivalent to the following ACL policy file entry:

```
entry udpacl {
  if {
    source-address 10.203.134.0/24;
    destination-address 140.158.18.16/32;
    protocol udp;
    source-port 190;
    destination-port 1200 - 1250;
  } then {
    permit;
  }
}
```

History

This command was first available in ExtremeXOS 11.3.

The **non_permanent** option was added in ExtremeXOS 11.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

create access-list network-zone

```
create access-list network-zone zone_name
```

Description

Creates a network-zone with a specified name.

Syntax Description

access-list	Access list
network-zone	Network zone
<i>zone_name</i>	Network zone name

Default

N/A.

Usage Guidelines

Use this command to create a network-zone with a specified name. The network-zone can then be associated with the policy file using either the "source-zone" or "destination-zone" attribute.

Example

```
Switch# create access-list network-zone zone1
```

If the user tries to create a network-zone that was already created, the following error message will be displayed on the console, and the command will be rejected.

```
Switch#create access-list network-zone zone1
Error: Network Zone "zone1" already exists.
```

History

This command was first available in ExtremeXOS 15.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

create access-list zone

```
create access-list zone name zone-priority number
```

Description

Creates a dynamic *ACL* zone, and sets the priority of the zone.

Syntax Description

<i>name</i>	Specifies the dynamic ACL zone name. The name can be from 1-32 characters long.
zone-priority <i>number</i>	Specifies priority of the zone. The range is from 1 (highest priority) to 4294967295 (lowest priority).

Default

The denial of service, system, and security zones are configured by default, and cannot be deleted.

Usage Guidelines

This command creates a dynamic ACL zone. You can configure the priority of the zone in relation to the default zones or to other configured zones.

Example

The following command creates a new zone, called myzone, with a priority of 2:

```
create access-list myzone zone-priority 2
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

create account

```
create account [admin | user | lawful-intercept] account-name {encrypted  
  encrypted_password | password}
```

Description

Creates a new user account.

Syntax Description

admin	Specifies an access level for admin account type. This user has read and write privileges.
user	Specifies an access level for user account type. This user has read-only privileges.
lawful-intercept	Specifies an access level for lawful intercept account type.
<i>account-name</i>	Specifies a new user account name.
encrypted	<p>Caution: Using this option incorrectly can result in you being locked out of your switch account.</p> <p>This option specifies that the entered password is in encrypted hash format, not that the resulting password will be stored in encrypted form. Generally, this option should not be used. Using this option with a plain text password, as opposed to a hashed version of a password, can result in the user being locked out of the account.</p>
<i>password</i>	Specifies a user password.

Default

N/A.

User Account Levels

By default, the switch is configured with two accounts with the access levels shown in the table below.

Account Name	Access Level
admin	You can access and change all manageable parameters. The admin account cannot be deleted.
user	<p>You can view (but not change) all manageable parameters, with the following exceptions:</p> <ul style="list-style-type: none"> You cannot view the user account database. You cannot view the <i>SNMP</i> community strings. You cannot view SSL settings. <p>This user has access to the ping command.</p>
lawful-intercept	<p>This user has special lawful intercept and read-only privileges.</p> <p>Note: Only a single lawful-intercept account can exist at any one time on the system.</p>

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. For name creation guidelines and a list of reserved names, see *Object Names* in the [Switch Engine 32.3 User Guide](#).

Usage Guidelines

The switch can have a total of 16 user accounts.

The system must have one administrator account.

When you use the **encrypted** keyword, the following password that you specify should be in encrypted hash format. Administrators should **not** use the encrypted option and should enter the password in plain text. Using this option with a plain text password, as opposed to a hashed version of a password, can result in the user being locked out of the account. Generally, this option should not be used. A valid use of this option would be when transferring account information between switches using the output of the [show configuration](#) on page 2569 command, where the displayed password is in hashed form. You can copy this hashed password and enter it as the password with the **encrypted** option. The switch will de-crypt the hashed password into the plain text password that as specified for the original account.

The system prompts you to specify a password after you enter this command and to reenter the password. If you do not want a password associated with the specified account, press **[Enter]** twice.

You must have administrator privileges to change passwords for accounts other than your own. User names are not case-sensitive. Passwords are case-sensitive. 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 32 characters. For user names, only alphanumeric, dash (-), and underscore (_) characters may be used. If you use a hashtag (#), everything after it is ignored.

**Note**

User names cannot begin with a number.

**Note**

If the account is configured to require a specific password format, the minimum is eight characters. See [configure account password-policy char-validation](#) for more information.

Example

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

```
create account admin John2
```

History

This command was first available in ExtremeXOS 10.1.

The **encrypted** option was added in ExtremeXOS 11.5.

The **lawful intercept** option was added in ExtremeXOS 15.3.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

create auto-peering bgp

```
create auto-peering bgp routerid ipaddress AS-number asNumber
```

Description

This command creates and enables BGP auto-peering using specified *BGP* router ID and AS number.

Syntax Description

routerid	Designates a BGP router ID.
<i>ipaddress</i>	Specifies the BGP router ID as an IP address in IPv4 format (x.y.z.w).
AS-number	Designates unique BGP Autonomous System (AS) number.
<i>asNumber</i>	Specifies the AS number (1–4,294,967,295).

Default

N/A

Usage Guidelines

This command creates VLANs dynamically. It also creates a loopback VLAN with an IP address of the BGP router ID. Within BGP, the router ID, AS number, and easyBGP capability are configured along with redistribution of host-mobility routes. Dynamic VLANs are created if no VLANs are specified.

A save and reboot is required if ECMP exceeds 16.

To view BGP auto-peering status, use the command `show auto-peering {bgp | ospf}`.

Example

The following example creates auto-peering using BGP router ID at 10.3.4.2 with AS 52:

```
# create auto-peering bgp routerid 10.3.4.2 AS-number 52
```

History

This command was first available in ExtremeXOS 22.5.

The requirement to specify a VLAN range was made optional in ExtremeXOS 30.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

This feature requires the Advanced Edge license. For more information about licenses, see the [Switch Engine 32.3 Feature License Requirements](#).

create auto-peering ospf

```
create auto-peering ospf routerid ipaddress
```

Description

Creates and enables OSPFv2 auto-peering.

Syntax Description

auto-peering	Specifies configuring Auto-peering.
ospf	Specifies configuring OSPF Auto-peering.
routerid	Specifies providing an OSPF router ID.
<i>ipaddress</i>	Sets the OSPF router IP address in IPv4 format (x.y.z.w).

Default

N/A.

Usage Guidelines

This command creates a loopback VLAN with the IP address of the supplied OSPF router ID. Within OSPF, the router ID, VXLAN-extensions are configured along with redistribution of host-mobility routes.

ECMP between two switches is not supported with OSPFv2 auto-peering. Only one link forms an adjacency, and traffic is lost on failover. Link aggregation is the preferred configuration for the topology.

To view OSPFv2 auto-peering status, use the command `show auto-peering {bgp | ospf}`.

Example

The following example creates and enables OSPFv2 auto-peering using with the OSPF router ID set to "10.3.4.2":

```
# create auto-peering ospf routerid 10.3.4.2
```

History

This command was first available in ExtremeXOS 30.6.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the OSPFv2 Auto-peering feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

create bgp evpn instance

```
create bgp evpn instance evpn_instance_name
```

Description

Creates an EVPN instance.

Syntax Description

bgp	BGP capability.
evpn	EVPN protocol.
instance	Specifies creating an EVPN instance
<i>evpn_instance_name</i>	Name of the EVPN instance.

Default

N/A.

Usage Guidelines

The EVPN instance will become active if the configured VNI matches the configured VNI of a virtual network.

Example

The following example creates an EVPN instance named "my_evpn":

```
# create bgp evpn instance my_evpn
```

History

This command was first available in ExtremeXOS 30.2.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

create bgp neighbor peer-group

```
create bgp neighbor remoteaddr peer-group peer-group-name {multi-hop}
```

Description

Creates a new neighbor and makes it part of the peer group.

Syntax Description

<i>remoteaddr</i>	Specifies the IPv4 or IPv6 address of a <i>BGP</i> neighbor.
<i>peer-group-name</i>	Specifies a peer group.
multi-hop	Specifies to allow connections to EBGp peers that are not directly connected.

Default

N/A.

Usage Guidelines

You can specify an IPv4 or IPv6 address for the BGP peer. The address can be a global unicast or a link-local address. IPv6 link-local remote addresses are supported only for EBGp single-hop peerings.

If you are adding an IPv4 peer to a peer group and no IPv4 address family capabilities are assigned to the specified peer group, the IPv4 unicast and multicast address families are automatically enabled for that peer group. If you adding an IPv6 peer to a peer group and no IPv6 address family capabilities are assigned to the peer group, you must explicitly enable the IPv6 address family capabilities you want to support.



Note

If the peer group or any member of the peer group has been configured with an IPv4 or IPv6 address family, the peer group only accepts peers that are configured to use that family. For example, if a peer group is configured for the IPv4 unicast address family, the switch will not allow you to add an IPv6 peer. Likewise, an IPv6 peer group cannot accept an IPv4 peer.

If the multihop keyword is not specified, the IP addresses of the EBGp speaker and peer must belong to the same subnet.

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:

```
configure bgp neighbor [all | remoteaddr] peer-group [peer-group-name | none] {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 outer:

```
create bgp neighbor 192.1.1.22 peer-group outer
```

The following example specifies how to create a neighbor peer group in a VRF (PE – CE neighbor session):

```
virtual-router <vr_vrf_name>
create bgp neighbor <remoteaddr> remote-AS-number <asNumber> {multi-hop}
create bgp neighbor <remoteaddr> peer-group <peer-group-name> {multi-hop}
delete bgp [{neighbor} <remoteaddr> | neighbor all ]
[create | delete] bgp peer-group <peer-group-name>
```

BGP maintains a separate RIB (RIB-In, RIB-Loc and RIB-Out) for each of the VRF it is configured to run. So routes received from a peer in VRF1 are not mixed up with routes from a peer in VRF2. Additionally, BGP routes in a VRF are regular IPv4 routes of address family ipv4. The BGP decision algorithm occurs inside a VRF and is not impacted by any BGP activity in other VRF. There can be two BGP neighbors with the same peer IP address in two different VRFs.

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Support for L3 VPN was added in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

create bgp neighbor remote-AS-number

```
create bgp neighbor remoteaddr remote-AS-number as-number {multi-hop}
```

Description

Creates a new BGP peer.

Syntax Description

<i>remoteaddr</i>	Specifies an IP address of the BGP neighbor.
<i>as-number</i>	Specifies a remote AS number. The range is 1 to 4294967295.
multi-hop	Specifies to allow connections to EBGp peers that are not directly connected.

Default

N/A.

Usage Guidelines

You can specify an IPv4 or IPv6 address for the BGP peer. The address can be a global unicast or a link-local address. IPv6 link-local remote addresses are supported only for EBGP single-hop peerings.

If the multihop keyword is not specified, the IP addresses of the EBGP speaker and peer must belong to the same subnet.

The AS number is a 4-byte AS number in either the ASPLAIN or the ASDOT format as described in RFC 5396, Textual Representation of Autonomous System (AS) Numbers.

If the AS number is the same as the AS number provided in the `configure bgp as` 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 specifies a BGP peer AS number using the ASPLAIN 4-byte AS number format:

```
create bgp neighbor 10.0.0.1 remote-AS-number 65540
```

The following command specifies a BGP peer AS number using the ASDOT 4-byte AS number format:

```
create bgp neighbor 10.0.0.1 remote-AS-number 1.5
```

The following command specifies a BGP peer using an IPv6 address:

```
create bgp neighbor fe80::204:96ff:fe1e:a8f1%vlan1 remote-AS-number 200
```

The following example specifies how to create a neighbor peer group in a VRF (PE – CE neighbor session):

```
virtual-router <vr_vrf_name>  
create bgp neighbor <remoteaddr> remote-AS-number <asNumber> {multi-hop}  
create bgp neighbor <remoteaddr> peer-group <peer-group-name> {multi-hop}  
delete bgp [{neighbor} <remoteaddr> | neighbor all ]  
[create | delete] bgp peer-group <peer-group-name>
```

BGP maintains a separate RIB (RIB-In, RIB-Loc and RIB-Out) for each of the VRF it is configured to run. So routes received from a peer in VRF1 are not mixed up with routes from a peer in VRF2. Additionally, BGP routes in a VRF are regular IPv4 routes of address family ipv4. The BGP decision algorithm occurs inside a VRF and is not impacted by any BGP activity in other VRF. There can be two BGP neighbors with the same peer IP address in two different VRFs.

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for 4-byte AS numbers was first available in ExtremeXOS 12.4.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Support for L3 VPN was added in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

create bgp peer-group

```
create bgp peer-group peer-group-name
```

Description

Creates a new peer group.

Syntax Description

<i>peer-group-name</i>	Specifies a peer group.
------------------------	-------------------------

Default

N/A.

Usage Guidelines

You can use BGP peer groups to group together up to 512 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:

- source-interface
- out-nlri-filter
- out-aspath-filter
- out-route-policy
- send-community
- next-hop-self

The BGP peer group name must begin with an alphabetical character and may contain alphanumeric characters and underscores (_), but it cannot contain spaces. The maximum allowed length for a name is 32 characters. For name creation guidelines and a list of reserved names, see the [Switch Engine 32.3 Feature License Requirements](#) document..

No IPv4 or IPv6 address family capabilities are added to a new peer group. When the first IPv4 peer is added to a peer group, the IPv4 unicast and multicast families are enabled by default. No IPv6 address family capabilities are automatically added when an IPv6 peer is added to a peer group; you must explicitly add any IPv6 address family capabilities that you want for a peer group.

Example

The following command creates a new peer group named outer:

```
create bgp peer-group outer
```

The following example specifies how to create a neighbor peer group in a VRF (PE – CE neighbor session):

```
virtual-router <vr_vrf_name>
create bgp neighbor <remoteaddr> remote-AS-number <asNumber> {multi-hop}
create bgp neighbor <remoteaddr> peer-group <peer-group-name> {multi-hop}
delete bgp [{neighbor} <remoteaddr> | neighbor all ]
[create | delete] bgp peer-group <peer-group-name>
```

BGP maintains a separate RIB (RIB-In, RIB-Loc and RIB-Out) for each of the VRF it is configured to run. So routes received from a peer in VRF1 are not mixed up with routes from a peer in VRF2. Additionally, BGP routes in a VRF are regular IPv4 routes of address family ipv4. The BGP decision algorithm occurs inside a VRF and is not impacted by any BGP activity in other VRF. There can be two BGP neighbors with the same peer IP address in two different VRFs.

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Support for L3 VPN was added in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

create cfm domain dns md-level

```
create cfm domain dns name md-level level
```

Description

Creates a maintenance domain (MD) in the DNS name format and assigns an MD level to that domain.

Syntax Description

<i>name</i>	Assigns the name you want for this domain, using the DNS name format. Enter alphanumeric characters for this format; the maximum is 43 characters.
<i>level</i>	Specifies the MD level you are assigning to this domain. Enter a value between 0 and 7.

Default

N/A.

Usage Guidelines

You can have up to 8 domains on a switch, and each one must have a unique MD level.

You assign each domain a maintenance domain (MD) level, which function in a hierarchy for forwarding CFM messages. The levels are from 0 to 7; with the highest number being superior in the hierarchy.

The IEEE standard 801.2ag specifies different levels for different network users, as follows:

- 5 to 7 for end users
- 3 and 4 for Internet service providers (ISPs)
- 0 to 2 for operators (entities carrying the information for the ISPs)



Note

MEPs with intervals 3 and 10 cannot be created in this domain as the domain name format is of dns type.

Example

The following command creates a domain, using the DNS name format, named extreme and assigns that domain an MD level of 2:

```
create cfm domain dns extreme md-level 2
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

create cfm domain mac md-level

```
create cfm domain mac mac-addr int md-level level
```

Description

Creates a maintenance domain (MD) in the MAC address + 2-octet integer format and assigns an MD level to that domain.

Syntax Description

<i>mac-addr</i>	Enter a MAC address in the format XX:XX:XX:XX:XX:XX to specify part of the domain name.
<i>int</i>	Enter the 2-octet integer you want to append to the MAC address to specify the domain name.
<i>level</i>	Specifies the MD level you are assigning to this domain. Enter a value between 0 and 7.

Default

N/A.

Usage Guidelines

You can have up to 8 domains on a switch, and each one must have a unique MD level.

You assign each domain a maintenance domain (MD) level, which function in a hierarchy for forwarding CFM messages. The levels are from 0 to 7; with the highest number being superior in the hierarchy.

The IEEE standard 801.2ag specifies different levels for different network users, as follows:

- 5 to 7 for end users
- 3 and 4 for Internet service providers (ISPs)
- 0 to 2 for operators (entities carrying the information for the ISPs)

Example

The following command creates a domain, using the MAC + 2-octet integer format, with the MAC address of 11:22:33:44:55:66 and an integer value of 63; it also assigns that domain an MD level of 2:

```
create cfm domain mac 11:22:33:44:55:66 63 md-level 2
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

create cfm domain string md-level

```
create cfm domain string str_name md-level level
```

Description

Creates a maintenance domain (MD) in the string name format and assigns an MD level to that domain.

Syntax Description

<i>str_name</i>	Enter a character string to specify part of the domain name. The maximum length is 43 characters.
<i>level</i>	Specifies the MD level you are assigning to this domain. Enter a value between 0 and 7.

Default

N/A.

Usage Guidelines

You can have up to 8 domains on a switch, and each one must have a unique MD level.

You assign each domain a maintenance domain (MD) level, which function in a hierarchy for forwarding CFM messages. The levels are from 0 to 7; with the highest number being superior in the hierarchy.

The IEEE standard 801.2ag specifies different levels for different network users, as follows:

- 5 to 7 for end users
- 3 and 4 for Internet service providers (ISPs)
- 0 to 2 for operators (entities carrying the information for the ISPs)

Example

The following command creates a domain, using the string format having a value of extreme; it also assigns that domain an MD level of 2:

```
create cfm domain string extreme md-level 2
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

create cfm segment destination

```
create cfm segment segment_name destination mac_addr {copy  
    segment_name_to_copy}
```

Description

Creates a CFM segment.

Syntax Description

<i>segment_name</i>	An alpha numeric string identifying the segment name.
<i>mac_addr</i>	Specifies the MAC address.
<i>segment_name_to_copy</i>	Specifies the CFM segment whose configuration is to be copied.

Default

N/A.

Usage Guidelines

Use this command to explicitly create a CFM segment where the segment name is a 32-byte long alpha-numeric character string.

Example

The following command creates a CFM segment named segment-new using MAC address 00:11:22:11:33:11 and copying segment-old:

```
create cfm segment segment-new destination 00:11:22:11:33:11 copy segment-old
```

Here, the *copy existing cfm segment* is an optional parameter, and if used, the following configurations from the existing CFM segment are copied to the newly created segment:

- DMM transmission interval
- Class of service
- Threshold values
- Measurement window size
- Timeout value



Note

The copy option is not shown in "show config" as it is used only for copying the existing values when creating a segment.

If you later configure any of the above mentioned information in segment-new, the old value(s) which were copied from segment-old will be overwritten with the new one in segment-new, as is done for any other commands. The same will not be true on the reverse case. If you modify the values of segment-old, the modified value will NOT be propagated to the CFM segments which use segment-old's configurations. In other words, the configurations of segment-old that are at the time of creating segment-new will alone be copied and not any other changes that are made to segment-old later on.

History

This command was first available in ExtremeXOS 12.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

create database

```
create database database_name
```

Description

Creates an Automation Edge remote VXLAN network identifier (VNI)-device database.

Syntax Description

database	Creates a remote VNI-database.
<i>database_name</i>	Specifies the name of the new database.

Default

N/A.

Usage Guidelines

You can only create one VNI-device database.

Example

The following example creates a database called "database1":

```
# create database database1
```

History

This command was first available in ExtremeXOS 31.1 as a demonstration feature.

This command is fully supported in ExtremeXOS 31.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

create eaps shared-port

```
create eaps shared-port ports
```

Description

Creates an EAPS shared port on the switch.

Syntax Description

<i>ports</i>	Specifies the port number of the common link port.
--------------	--

Default

N/A.

Usage Guidelines

To configure a common link, you must create a shared port on each switch on either end of the common link.

Example

The following command creates a shared port on the EAPS domain.

```
create eaps shared-port 1:2
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms with the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and what licenses are appropriate for this feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

create eaps

```
create eaps name
```

Description

Creates an EAPS domain with the specified name.

Syntax Description

<i>name</i>	Specifies the name of an EAPS domain to be created. Can be up to 32 characters in length.
-------------	---

Default

N/A.

Usage Guidelines

An EAPS domain name must begin with an alphabetical character and may contain alphanumeric characters and underscores (_), but it cannot contain spaces. The maximum allowed length for a name is 32 characters. For name creation guidelines and a list of reserved names, see [Object Names](#) in the ExtremeXOS Concepts Guide.

Example

The following command creates EAPS domain eaps_1:

```
create eaps eaps_1
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

create erps ring

```
create erps ring-name {ring-id ring_id}
```

Description

Creates an ERPS ring.

Syntax Description

<i>ring-name</i>	Alphanumeric string that identifies the ERPS ring.
ring-id	Specifies configuring a unique integer ID for ERPS ring.
<i>ring_id</i>	Sets the ERPS ring ID value. Range is 1 to 239.

Default

N/A.

Usage Guidelines

Use this command to create an ERPS ring, and optionally the ring ID.

Example

The following command creates an ERPS ring named "ring1" with ring ID "50":

```
create erps ring1 ring-id 50
```

History

This command was first available in ExtremeXOS 15.1.

Ring ID was added in ExtremeXOS 32.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

create esrp

```
create esrp esrp_domain {type [vpls-redundancy | standard]}
```

Description

Creates an ESRP domain with the specified name on the switch.

Syntax Description

<i>esrp_domain</i>	Specifies the name of an ESRP domain to be created. Can be up to 32 characters in length.
--------------------	---

Default

The ESRP domain is disabled and in the “Aware” state.

When you create an ESRP domain, it has the following default parameters:

- Operational version—Extended
- Priority—0
- VLAN interface—none
- VLAN tag—0
- Hello timer—2 seconds
- Neighbor timer—8 seconds
- Premaster timer—6 seconds
- Neutral timer—4 seconds
- Neighbor restart timer—30 seconds
- VLAN tracking—none
- Ping tracking—none
- IP route tracking—none

Usage Guidelines

The type keyword specifies the type of ESRP domain when a new ESRP domain is created. The only types supported are vpls-redundancy and standard. Not specifying the optional ESRP domain type results in the creation of an ESRP domain of type standard. The standard ESRP domain is equivalent

to the legacy ESRP domain type that was implicitly created. The vpls-redundancy domain type is only specified when redundant access to an *MPLS* VPLS network is desired.

An ESRP domain name must begin with an alphabetical character and may contain alphanumeric characters and underscores (_), but it cannot contain spaces. The maximum allowed length for a name is 32 characters. For ESRP domain name guidelines and a list of reserved names, see Object Names in the [Switch Engine 32.3 User Guide](#).

Each ESRP domain name must be unique and cannot duplicate any other named ESRP domains on the switch. If you are uncertain about the ESRP names on the switch, use the `show esrp` command to view the ESRP domain names.

You can create a maximum of 128 ESRP domains.

Configuring ESRP-Aware Switches

For an Extreme Networks switch to be ESRP-aware, you must create an ESRP domain on the aware switch, add a master VLAN to that ESRP domain, add a member VLAN to that ESRP domain if configured, and configure a domain ID if necessary.

For complete information about software licensing, including how to obtain and upgrade your license and what licenses are appropriate for this feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

Example

The following command creates ESRP domain esrp1 on the switch:

```
create esrp esrp1
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

create fdb mac-tracking entry

```
create fdb mac-tracking entry mac_addr
```

Description

Adds a MAC address to the MAC address tracking table.

Syntax Description

<i>mac_addr</i>	Specifies a device MAC address, using colon-separated bytes.
-----------------	--

Default

The MAC address tracking table is empty.

Usage Guidelines

None.

Example

The following command adds a MAC address to the MAC address tracking table:

```
create fdb mac-tracking entry 00:E0:2B:12:34:56
```

History

This command was first available in ExtremeXOS 12.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

create fdb vlan ports

```
create fdb mac_addr vlan vlan_name [ports port_list {tagged tag} |
blackhole | vxlan { vr vr_name } {ipaddress} remote_ipaddress ]
| broadcast vlan vlan_name vxlan { vr vr_name } {ipaddress}
remote_ipaddress | unknown-multicast vlan vlan_name vxlan { vr
vr_name } {ipaddress} remote_ipaddress | unknown-unicast vlan
vlan_name vxlan { vr vr_name } {ipaddress} remote_ipaddress ]
```

Description

Creates a permanent static *FDB* entry.

Syntax Description

<i>mac_addr</i>	Specifies a device MAC address, using colon-separated bytes.
<i>vlan_name</i>	Specifies a <i>VLAN</i> name associated with a MAC address.
<i>port_list</i>	Specifies one or more ports or slots and ports associated with the MAC address.
tagged tag	Specifies the port-specific VLAN tag. When there are multiple ports specified in <i>port_list</i> , the same tag is used for all of them.

blackhole	Enables the blackhole option. Any packets with either a source MAC address or a destination MAC address matching the FDB entry are dropped.
broadcast	Forwarding destination(s) for broadcast traffic.
unknown-unicast	Forwarding destination(s) for unknown unicast traffic.
<i>unknown-multicast</i>	Forwarding destination(s) for unknown multicast traffic.
vxlان	The MAC address is reachable through a VXLAN Tunnel.
vr	VR/VRF instance the IPv4 address is configured on.
<i>vr_name</i>	An existing VR/VRF name.
ipaddress	Configure the IP address of the remote tunnel endpoint to which the MAC needs to be bound.
<i>remote_ipaddress</i>	IPv4 address of the remote tunnel endpoint.

Default

N/A.

Usage Guidelines

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. After they have been created, permanent static entries stay the same as when they were created. If the same MAC address and VLAN 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 identifier (VLANid) is changed.
- A port is disabled.
- A port enters blocking state.
- A port goes down (link down).

A permanent static FDB entry is deleted when any of the following take place:

- A VLAN is deleted.
- A port mode is changed (tagged/untagged).
- A port is deleted from a VLAN.

Permanent static entries are designated by `spm` in the flags field of the `show fdb` output. You can use the `show fdb` command to display permanent FDB entries.

If the static entry is for a PVLAN VLAN that requires more than one underlying entry, the system automatically adds the required entries. For example, if the static entry is for a PVLAN network VLAN, the system automatically adds all required extra entries for the subscriber VLANs.

You can create FDB entries to multicast MAC addresses and list one or more ports. If more than one port number is associated with a permanent MAC entry, packets are multicast to the multiple destinations.

IGMP snooping rules take precedence over static multicast MAC addresses in the IP multicast range (01:00:5e:xx:xx:xx) unless IGMP snooping is disabled.

**Note**

When a multiport list is assigned to a unicast MAC address, load sharing is not supported on the ports in the multiport list.

In ExtremeXOS 21.1, this command was extended to add a remote VTEP as a destination to a MAC address. Three new tokens “broadcast”, “unknown-multicast” and “unknown-unicast” have been added to this command. When you want to specify a destination to forward all broadcast or unknown unicast traffic on that VLAN, these token are used. For “broadcast”, “unknown-multicast” and “unknown-unicast” only remote VTEPs (and not port_list or blackhole) can be specified in this release of ExtremeXOS. These entries can only be created when the virtual-network is in explicit-remote flooding mode.

Example

The following command adds a permanent, static entry to the FDB for MAC address 00 E0 2B 12 34 56, in VLAN marketing on port 4 on a switch:

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

The following example adds a permanent, static entry to the FDB for MAC address 00:01:02:03:04:05, in VLAN marketing, on a VLAN port that has tag 100 on port 3 on a switch:

```
create fdb 00:01:02:03:04:05 vlan msk ports 3 tag 100
```

History

This command was first available in ExtremeXOS 10.1.

The ability to create a multicast FDB with multiple entry ports was added in ExtremeXOS 11.3.

The **blackhole** option was first available for all platforms in ExtremeXOS 12.1.

In ExtremeXOS 12.3, the **fdb** keyword was introduced as an alias to the **fdbentry** keyword to avoid interference with the syntax of the MAC-Tracking feature commands. Both keywords execute; however, the syntax helper (tab completion) does not recognize the **fdbentry** keyword.

The **tag** keyword and example was added in ExtremeXOS 15.4.

Three new tokens “broadcast”, “unknown-multicast” and “unknown-unicast” were added to this command in ExtremeXOS 21.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

create flow-redirect

```
create flow-redirect flow_redirect_name
```

Description

Creates a named flow redirection policy.

Syntax Description

<i>flow_redirect_name</i>	Specifies the name of the flow redirection policy.
---------------------------	--

Default

N/A.

Usage Guidelines

Use this command to create a named flow redirection policy to which nexthop information can be added.

For name creation guidelines and a list of reserved names, see *Object Names* in the [Switch Engine 32.3 User Guide](#).

Example

The following example creates a flow redirection policy names flow3:

```
create flow-redirect flow3
```

History

This command was first available in ExtremeXOS 12.1.

The maximum number of flow redirects was increased to 4096 in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

create flowmon collector

```
create flowmon collector collector_name
```

Description

Creates a collector where Flow Monitor sends information.

Syntax Description

collector	Specifies to send flow information to a collector.
<i>collector_name</i>	Specifies the name of the created collector. Range is 32 characters.

Default

N/A.

Usage Guidelines

Up to eight external collectors are supported. The created collector must be configured before it can be added to a group, and a collector can be used by many groups.

The system will reject any attempt to create a collector that already exists.

Example

The following command creates a collector with the name 'src-ipv4-address':

```
# create flowmon collector src-ipv4-address
```

History

This command was first available in ExtremeXOS 32.2.

Platform Availability

This command is available on ExtremeSwitching 5520 and 5420 series switches.

create flowmon group

```
create flowmon group group_name
```

Description

Creates a Flow Monitor group.

Syntax Description

group	Specifies the Flow Monitor group.
<i>group_name</i>	Specifies the assigned name of the Flow Monitor group. Range is 32 characters.

Default

N/A.

Usage Guidelines

The system will reject any attempt to create a group that already exists.

Example

The following command creates a group with the name 'max-flow-age':

```
# create flowmon group max-flow-age
```

History

This command was first available in ExtremeXOS 32.2.

Platform Availability

This command is available on ExtremeSwitching 5520 and 5420 series switches.

create flowmon key

```
create flowmon key key_name
```

Description

Creates a Flow Monitor key.

Syntax Description

key	Specifies the Flow Monitor key.
<i>key_name</i>	Specifies the assigned name of the Flow Monitor key. Range is 32 characters.

Default

N/A.

Usage Guidelines

The system will reject any attempt to create a key that already exists.

Example

The following command creates a key with the name 'src-ipv4-addr':

```
# create flowmon key src-ipv4-addr
```

History

This command was first available in ExtremeXOS 32.2.

Platform Availability

This command is available on ExtremeSwitching 5520 and 5420 series switches.

create identity-management role

```
create identity-management role role_name match-criteria match_criteria
    {priority pri_value}
```

Description

Creates and configures an identity management role.

Syntax Description

<i>role_name</i>	Specifies a name for the new role (up to 32 characters).
<i>match_criteria</i>	Specifies an expression that identifies the users to be assigned to the new role.
<i>pri_value</i>	Specifies the role priority; the lower the priority number, the higher the priority. The range of values is 1 to 255. Value 1 represents the highest priority, and value 255 represents the lowest priority.

Default

Priority=255.

Usage Guidelines

The identity management feature supports a maximum of 64 roles.

The role name can include up to 32 characters. Role names must begin with an alphabetical letter, and only alphanumeric, underscore (_), and hyphen (-) characters are allowed in the remainder of the name. Role names cannot match reserved keywords, or the default role names reserved by identity manager. For more information on role name requirements and a list of reserved keywords, see *Object Names* in the [Switch Engine 32.3 User Guide](#). The role names reserved by identity manager are:

- authenticated.
- blacklist.
- unauthenticated.
- whitelist.

The match-criteria is an expression or group of expressions consisting of identity attributes, operators and attribute values. The maximum number of attribute value pairs in a role match criteria is 16. The variables in the match criteria can be matched to attributes retrieved for the identity from an LDAP server, or they can be matched to attributes learned locally by identity manager.

[Table 20](#) lists match criteria attributes that can be retrieved from an LDAP server.

[Table 21](#) on page 1624 lists locally learned attributes that can be used for match criteria.

Table 22 on page 1625 lists the match criteria operators.

Table 20: LDAP Match Criteria Attributes

LDAP Attribute Name	Value Type
l or location	String
company	String
co or country	String
department	String
employeeID	String
st or state	String
title	String
mail or email	String
memberOf	String

Table 21: Locally Learned Match Criteria Attributes

Attribute Description	Attribute Name	Value Type	Example
LLDP device name	device-model	String	device-name == Avaya4300
LLDP device capabilities	device-capability	String:OtherRepeaterBridgeWLAN access portRouterPhoneDOCSIS cable deviceStation only	device- capability == Telephone
LLDP device manufacturer name	device-manufacturer-name	String	device- manufacturer- name == Avaya
LLPD system description	device-description	String	device- description==Del l EqualLogic Storage Array
MAC address	mac	MAC	mac == 00:01:e6:00:00:0 0/ ff:ff:ff:00:00:0 0
MAC OUI	mac-oui	MAC	mac-oui == 00:04:96
IP address	ip-address	IP	ip-address == 10.1.1.0/20

Table 21: Locally Learned Match Criteria Attributes (continued)

Attribute Description	Attribute Name	Value Type	Example
User name	username	String	userName == adam
Port list	ports	Portlist	ports == 1,5-8

Table 22: Match Criteria Operators

Operator	Description
==	Equal. Creates a match when the value returned for the specified attribute matches the value specified in the role.
!=	Not equal. Creates a match when the value returned for the specified attribute does not match the value specified in the role.
AND	And. Creates a match when the two expressions joined by this operator are both true.
contains	Contains. Creates a match when the specified attribute contains the text specified in the role definition.
;	Semicolon. This delimiter separates expressions within the match criteria.

The role priority determines which role a user is mapped to when the user's attributes match the match-criteria of more than 1 role. If the user's attributes match multiple roles, the highest priority (lowest numerical value) role applies. If the priority is the same for all matching roles, the role for which the priority was most recently set or modified is used.

Example

The following examples create roles for the conditions described in the comments that precede the commands:

```
# Creates a role named "India-Engr" that matches employees from the Engineering
# department who work in India
* Switch.22 # create identity-management role "India-Engr" match-criteria
"country==India; AND department==Engineering;"
# Creates a role named "US-Engr" that matches employees whose title is Engineer and
# who work in United States
* Switch.23 # create identity-management role US-Engr match-criteria "title contains
Engineer; AND country == US;" priority 100
# Creates a role named "Avaya4300Device" for Avaya phones of type 4300 that are
# manufactured by Avaya
* Switch.24 # create identity-management role "Avaya4300Device" match-criteria "device-
capability == Phone; AND device-name == Avaya4300; AND device-manufacturer-name == Avaya;"
# Creates a role for all Extreme Networks switches with MAC-OUI "00:04:96"
* Switch.25 # create identity-management role "ExtremeSwitch" match-criteria "mac-oui ==
00:04:96;"
# Creates a role for all identities with IP address 1.2.3.1 - 1.2.3.255
* Switch.26 # create identity-management role "EngineeringDomain" match-criteria "ip-
Address == 1.2.3.0/255.255.255.0;"
# Creates a role for all phone devices with MAC_OUI of "00:01:e6"
* Switch.27 # create identity-management role "Printer" match-criteria "mac ==
00:01:e6:00:00:00/ff:ff:ff:00:00:00; device-capability == Phone;"
# Creates a role for the user name "adam" when he logs in from IP address 1.2.3.1 -
# 1.2.3.255.
```

```
* Switch.28 # create identity-management role "NotAccessibleUser" match-criteria
"userName == adam; AND "ip-Address == 1.2.3.0/24;"
# Creates a role named "secureAccess" for users who log in on ports 1, 5, 6, 7, and 8
# with IP addresses in the range of 10.1.1.1 to 10.1.1.255
create identity-management role "SecureAccess" match-criteria "ports == 1,5-8; AND ip-
address == 10.1.1.0/20;"
# Creates a role named "Prod-Engineers" for all the engineers who are under LDAP group
'Production'.
Create identity-management role "Prod-Engineers" match-criteria "title==Engineer; AND
memberOf==Production;"
```

History

This command was first available in ExtremeXOS 12.5.

Support for matching locally learned attributes was added in ExtremeXOS 12.7.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

create ip nat rule

```
create ip nat rule rule_name type [ source-nat | napt | destination-
napt]
```

Description

Creates an IP Network Address Translation (NAT) rule.

Syntax Description

ip	Specifies Internet Protocol (IP).
nat	Specifies NAT.
rule	Specifies creating a NAT rule.
<i>rule_name</i>	Specifies the NAT rule name.
type	Specifies the NAT translation type.
source-nat	Specifies modifying the source IP address for outbound traffic.
napt	Specifies modifying the source IP address and transport identifier for outbound traffic (Network Address Port Translation).
destination-napt	Specifies modifying the destination IP address and transport identifier for inbound traffic.

Default

N/A.

Usage Guidelines

The **type** option specifies the kind of translation that is carried out for the NAT rule.

To delete a rule, run the command `delete ip nat rule rule_name`.

User-created rules cannot be created with a name starting with "SYS_NAT_RULE_".

Example

The following example creates an IP NAT rule named "rule1" where the translation modifies the source IP address for outbound traffic:

```
# create ip nat rule rule1 type source-nat
```

History

This command was first available in ExtremeXOS 31.2.

Platform Availability

This command is available on the switches that support the IP NAT feature. For information about which switches support this and other features, see the [Switch Engine 32.3 Feature License Requirements](#) document.

create isis area

```
create isis area area_name
```

Description

This command creates an IS-IS router process in the current virtual router.

Syntax Description

<i>area_name</i>	Defines a name for the new IS-IS router process.
------------------	--

Default

N/A.

Usage Guidelines

No PDUs are sent until after the following events:

- The router process has been enabled
- The router process has been assigned a system ID and area address
- The router process has at least one interface (VLAN) that has IPv4 or IPv6 forwarding enabled.

By default, newly created IS-IS router processes are Level 1/Level 2 routers if a level 2 router process does not already exist in the current virtual router. No more than one IS-IS router process may be configured as a level 2 router. IS-IS router processes on different virtual routers may have the same name, but this is not recommended as it may cause confusion when administering the switch. The router process name supplied with this command may be optionally used as the hostname for this router process when dynamic hostname exchange support is enabled.

The area name must begin with an alphabetical character and may contain alphanumeric characters and underscores (_), but it cannot contain spaces. The maximum allowed length for a name is 32 characters. For name creation guidelines and a list of reserved names, see [Object Names](#) in the ExtremeXOS Concepts Guide.

A maximum of one area can be created per VR in this release.

Example

The following command creates a new IS-IS router process named areax:

```
create isis area areax
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Premier license as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

create keychain

```
create keychain keychain_name
```

Description

This command creates a keychain.

Syntax Description

<i>keychain_name</i>	Defines a name for the new keychain. The range is 1-31.
----------------------	---

Default

N/A.

Usage Guidelines

Use this command to create a keychain.

Example

The following command creates a new OSPFv3 keychain:

```
create keychain ospfv3-keys1
```

History

This command was first available in ExtremeXOS 31.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

create l2pt profile

```
create l2pt profile profile_name
```

Description

Creates an L2PT profile.

Syntax Description

l2pt	Creates a Layer 2 protocol tunneling profile.
profile	Profile that defines L2PT configuration for L2 protocols.
<i>profile_name</i>	Specifies a profile name (maximum 32 characters).

Default

Disabled.

Usage Guidelines

Use this command to create an L2PT profile.

Example

The following example create a new L2PT profile named "my_l2pt_prof":

```
create l2pt profile my_l2pt_prof
```

History

This command was first available in ExtremeXOS 15.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

create l2vpn fec-id-type pseudo-wire

```
create l2vpn [vppls vppls_name | vpws vpws_name] fec-id-type pseudo-wire
           pwid
```

Description

Creates a Layer 2 VPN, which can be either a VPLS or VPWS.

Syntax Description

<i>vppls_name</i>	Identifies the VPLS within the switch (character string). The <i>vppls_name</i> string must begin with an alphabetic character, and may contain up to 31 additional alphanumeric characters.
<i>vpws_name</i>	Identifies the VPWS within the switch (character string). The <i>vpws_name</i> string must begin with an alphabetic character, and may contain up to 31 additional alphanumeric characters.
<i>pwid</i>	Specifies a PW ID. Must be a non-zero 32-bit value that has network-wide significance.

Default

For the VPLS dot1q tag, the default value is exclude.

Usage Guidelines

Each VPLS or VPWS is a member of a single VPN, and each VPN can have only one associated VPLS or VPWS per switch. External to the switch, each VPN has an identifier.

A VPLS or VPWS name must begin with an alphabetical character and may contain alphanumeric characters and underscores (_), but it cannot contain spaces. The maximum allowed length for a name is 32 characters. For name creation guidelines and a list of reserved names, see *Object Names* in the [Switch Engine 32.3 User Guide](#).

Any non-zero 32-bit value that has network-wide significance can be specified for the identifier. This *pwid* is used on all pseudo-wires in the VPLS.

The **l2vpn** keyword is introduced in ExtremeXOS Release 12.4 and is required when creating a VPWS. For backward compatibility, the *l2vpn* keyword is optional when creating a VPLS. However, this keyword will be required in a future release, so we recommend that you use this keyword for new configurations and scripts.



Note

The switch's LSR ID must be configured before a VPLS or VPWS can be created.

Example

This example creates a VPLS with 99 as the PW ID:

```
create vpls vpls1 fec-id-type pseudo-wire 99
```

The following example creates a VPWS with 101 as the PW ID:

```
create l2vpn vpws vpws1 fec-id-type pseudo-wire 101
```

History

This command was first available in ExtremeXOS 11.6.

The l2vpn and vpws keywords were first available in ExtremeXOS 12.4.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

create ldap domain

```
create ldap domain domain_name {default}
```

Description

This command is used to add an LDAP domain. The new domain can be added as the default. Older default domains, if any, will no longer be the default since once only one domain can be default at a time.

Syntax Description

<i>domain_name</i>	Name of new LDAP domain to be added
--------------------	-------------------------------------

Default

N/A.

Usage Guidelines

Use this command to add an LDAP domain.

You can see the LDAP domains added by using the show ldap domain command.

Supporting multiple domains gives ExtremeXOS the capability to send LDAP queries to gather information about users belonging to different domains but connected to the same switch.

You can add upto 8 LDAP domains.

Example

The following command creates an LDAP domain with the name "sales.XYZCorp.com" and marks it as the default domain:

```
create ldap domain sales.XYZCorp.com default
```

History

This command was first available in ExtremeXOS 15.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

create log filter

```
create log filter name {copy filter_name}
```

Description

Creates a log filter with the specified name.

Syntax Description

<i>name</i>	Specifies the name of the filter to create.
copy	Specifies that the new filter is to be copied from an existing one.
<i>filter_name</i>	Specifies the existing filter to copy.

Default

N/A.

Usage Guidelines

This command creates a filter with the name specified. A filter is a customizable list of events to include or exclude, and optional parameter values. The list of events can be configured by component or subcomponent with optional severity, or individual condition, each with optional parameter values. See the commands [configure log filter events](#) and [configure log filter events match](#) for details on how to add items to the filter.

The filter can be associated with one or more targets using the [configure log target filter](#) command to control the messages sent to those targets. The system has one built-in filter named DefaultFilter, which itself may be customized. Therefore, the [create log filter](#) command can be used if a filter other than DefaultFilter is desired. As its name implies, DefaultFilter initially contains the default level of logging in which every ExtremeXOS component and subcomponent has a pre-assigned severity level.

If another filter needs to be created that will be similar to an existing filter, use the copy option to populate the new filter with the configuration of the existing filter. If the copy option is not specified, the new filter will have no events configured and therefore no incidents will pass through it.

The total number of supported filters, including DefaultFilter, is 20.

Example

The following command creates the filter named fdb2, copying its configuration from the filter DefaultFilter:

```
create log filter fdb2 copy DefaultFilter
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

create log message

```
create log message text
```

Description

This command logs an event using the text provided as the message.

Syntax Description

log	Configure log service.
message	Message to be logged.
<i>text</i>	Text of log message.

Default

N/A.

Usage Guidelines

Use this command to log an event using the *text* as provided as the message.

Example

```
# create log message "Creating the test VLAN"
# show log
08/06/2012 14:11:28.28 <Info:System.userComment> Creating the test VLAN
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

create log target upm

```
create log target upm {upm_profile_name}
```

Description

Creates a new UPM target profile.

Syntax Description

<i>upm_profile_name</i>	Specifies the name of an existing UPM profile.
-------------------------	--

Default

N/A.

Usage Guidelines

After configuration, a UPM log target links an EMS filter with a UPM profile. This command creates the UPM log target.

The default configuration for a new log target binds the target to the EMS filter defaultFilter, which is used for all system events. To configure the log target, use the command: `configure log target upm {upm_profile_name} filterfilter-name {severity [[severity] {only}]}`.

The default status of a new UPM log target is disabled. To enable the log target, use the command: `enable log target upm {upm_profile_name}`.

To view the log target, use the command: `show log configuration target upm {upm_profile_name}`.

Example

The following example creates a new UPM log target named testprofile1:

```
create log target upm testprofile1
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the Universal Port feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

create log target xml-notification

```
create log target xml-notification [ target_name | xml_target_name ]
```

Description

Creates a Web server XML-notification target name.

Syntax Description

<i>target_name</i>	Specifies the name of a non-existing XML notification target.
<i>xml_target_name</i>	Specifies the name of an already existing XML notification target.

Default

N/A.

Usage Guidelines

Use this command to create a web server XML-notification target name for EMS.

Example

The following command creates the target name test2:

```
create log target xml-notification test2
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

create macsec connectivity-association

```
create macsec connectivity-association ca_name pre-shared-key ckn ckn
cak [encrypted encrypted_cak | cak]
```

Description

Creates a named connectivity-association (CA) object that holds MAC Security (MACsec) key authentication data.

Syntax Description

connectivity-association	Secures connectivity provided between MACsec stations.
<i>ca_name</i>	Defines CA object name.
pre-shared-key	Selects static MACsec key consisting of both a CKN and CAK:
ckn	Selects CA key name. This public (non-secret) key name allows each of the MKA participants to select which connectivity association key (CAK) to use to process a received MACsec key agreement (MKA) protocol packets (MKPDU).
<i>ckn</i>	Sets the CA key name. Length allowed is 1–32 characters, entered as ASCII or an octet string preceded with 0x.
cak	Sets the connectivity association key (CAK). If you are using 256-bit cipher suite, then the CAK must be 32 octets. The 128-bit cipher suite can use either a 16- or 32-octet CAK. This is a long-lived secret key used to derive short-lived lower-layer keys (ICK, KEK, and SAK) which are used for key distribution and data encryption.
<i>cak</i>	Sets the non-encrypted CAK value. Must be entered as an octet string (for example: "0x859e72f0..."). A 128-bit (16 octet) CAK requires 32 hexadecimal digits, and a 256-bit (32 octet) CAK requires 64 hexadecimal digits. These values are secret and should be generated off switch with a suitable pseudorandom number generator.
encrypted	Designates that secret key value is in encrypted format.
<i>encrypted_cak</i>	Sets the value for the secret key. The encrypted CAK value is generated by the <code>show configuration macsec</code> command for previously configured CAKs.

Default

N/A.

Usage Guidelines

Up to 64 unique CA profiles can be created.

Example

The following example creates the CA object "testca" with a CKN of "the blue key" and 128-bit CAK of "0x01020304050607080910111213141516":

```
# create macsec connectivity-association testca pre-shared-key ckn "the blue key" cak
"0x01020304050607080910111213141516"
```

The following example creates the CA object "testca2" with a CKN of "the red key" and 256-bit CAK of "0x0102030405060708091011121314151617181920212223242526272829303132":

```
# create macsec connectivity-association testca2 pre-shared-key ckn "the red key" cak
"0x0102030405060708091011121314151617181920212223242526272829303132"

# show macsec connectivity-association
MACsec CAK Bit
CA Name Ports Length CAK Name (CKN)
-----
testca None 128 the blue key
testca2 None 256 the red key
```



Note
The CAKs shown here are examples. Use your own random number for maximum security.

History

This command was first available in ExtremeXOS 30.1.

Platform Availability

This command is available on the following platforms.



Note
The MACsec feature requires the installation of the MAC Security feature pack license.

Platform	Ports
ExtremeSwitching 5320	All ports of all models except stacking ports.
ExtremeSwitching 5420	All ports of all models except stacking ports.
ExtremeSwitching 5520	All ports, except 5520-VIM-4X and 5520-24X 10G ports
ExtremeSwitching 5720	All ports of all models except stacking ports.

create meter

```
create meter meter-name
```

Description

This command creates a meter for ingress traffic rate limiting.

Syntax Description

<i>meter-name</i>	Specifies the meter name.
-------------------	---------------------------

Default

N/A.

Usage Guidelines

Meter names must begin with an alphabetical character and may contain alphanumeric characters and underscores (_), but they cannot contain spaces. The maximum allowed length for a name is 32 characters. For meter name guidelines and a list of reserved names, see *Object Names* in the [Switch Engine 32.3 User Guide](#).

Example

The following command creates the meter maximum_bandwidth:

```
create meter maximum_bandwidth
```

History

This command was available in ExtremeXOS 11.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

create mirror control_index

```
create mirror control_index
```

Description

Creates a "control group" mirror referenced by a unique control index.

Syntax Description

<i>control_index</i>	Mirror destination control index in the form of a number 1–4. Also know as. etsysMirrorDestinationControlIndex. Each comprises a group of mirror names.
----------------------	---

Default

N/A.

Usage Guidelines

You can apply mirrors to policy profile rules by using a "control group" mirror referenced by a unique control index number (1-4). These control group mirrors are etsysMirrorDestinationControlEntry entries in the ENTERASYS-MIRROR-CONFIG-MIB (Mirror MIB). A Mirror MIB instance (designated by a control index) can be associated with up to four "physical" mirrors, each being one destination port (or tunnel). To create physical mirrors, use the command `create mirror mirror_name {to [port port | port-list port_list loopback-port port] { remote-tag rtag } | remote-ip remote_ip_address {{ vr } {vr_name } {from [source_ip_address | auto-source-ip] } {ping-check [on | off] } priority priority_value }} {description mirror-desc}`.

Example

The following example creates a control group mirror with control index number of "1":

```
# create mirror 1
```

History

This command was first available in ExtremeXOS 30.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

create mirror

```
create mirror mirror_name {to [port port | port-list port_list loopback-  
port port] { remote-tag rtag } | remote-ip remote_ip_address {{ vr }  
{vr_name } {from [ source_ip_address | auto-source-ip] } {ping-check  
[on | off] } priority priority_value }} {description mirror-desc}
```

Description

Creates a named mirror instance with an optional description, and optional "to port" definition, or remote IP address destination.

Syntax Description

<i>mirror_name</i>	Specifies the mirror name.
port	Specifies the mirror output port.
<i>port_list</i>	Specifies the list of ports where traffic is to be mirrored.

loopback-port	Specifies an otherwise unused port required when mirroring to a port_list. The loopback-port is not available for switching user data traffic.
<i>port</i>	Specifies a single loopback port that is used internally to provide this feature.
remote-tag	Specifies the value of the <u>VLAN</u> ID used by the mirrored packets when egressing the monitor port.
description	Specifies a description of the named mirror instance.
<i>mirror-desc</i>	The specified mirror description.
remote-ip	Specifies to send mirrored packets to specified remote destination IP address.
<i>remote_ip_address</i>	Specifies the destination remote IP address for mirrored packets.
vr	Specifies a virtual router of the remote IP address.
<i>vr_name</i>	Specifies the virtual router name. If not specified, VR of current command context is used.
from	Configures source IP address of encapsulated mirrored packets.
<i>source_ip_address</i>	Specifies the local source IPv4 address for encapsulated mirrored packets.
auto-source-ip	Automatically use source IP address of egress VLAN to be used to reach remote IP address.
ping-check	Configure ping health check for remote IP address.
on	Only send mirrored packets to remote IP address if periodic pings to remote IP address are successful (default).
off	Send mirrored packets to remote IP address without any ping health check, assuming MAC address and port of next hop IP address are static or learned.
priority	Configures a unique priority value for each redundant remote IP address of a mirror instance.
<i>priority_value</i>	Sets the unique priority value for the remote IP address. The priority value must be unique for each remote IP address in the mirror instance. The range is from 1 (least preferred) to 100 (most preferred). The default is 50.

Default

Disabled.

Ping health check of the remote IP address is enabled unless otherwise specified.

If a VR is not specified, the VR of the current command context is used.

The default priority value is 50.

Usage Guidelines

Use this command to create a named mirror instance with an optional description and optional "to port" or remote IP address definitions. You can create 15 named instances (the instance "DefaultMirror" is created automatically).

For high availability, you can add up to four redundant remote IP addresses. When creating a mirror with this command, you can add one IP address. To add additional remote IP addresses, use the configure mirror *mirror_name* {**to** [**port** *port* | **port-list** *port_list* | **loopback** **port** *port*] | **remote-ip** {**add**} *remote_ip_address* [{**vr**} *vr_name*] {**from** [*source_ip_address* | **auto-source-ip**] } {**ping-check** [**on** | **off**] } {**remote-tag** *rtag* | **port** **none**} {**priority** *priority_value*}command.

Example

The following example creates a mirror instance on port 3, slot 4 :

```
create mirror to port 3:4
```

History

This command was first available in ExtremeXOS 15.3.

The remote IP address option was added in ExtremeXOS 22.4.

Redundant remote IP addresses capability was added in ExtremeXOS 30.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

create mlag peer

```
create mlag peer peer_name { authentication [ md5 key {encrypted
    encrypted_auth_key | auth_key ]}] }
```

Description

Creates an peer switch association structure.

Syntax Description

<i>peer_name</i>	Specifies an alpha numeric string identifying the MLAG peer.
authentication	Authentication for MLAG checkpoint connection.
md5	<i>MD5</i> authentication type.
key	Authentication key for checkpoint connection to the MLAG peer.
encrypted	Authentication key in in encrypted format.
<i>auth_key</i>	Authetication key used for checkpoint connection.

Default

N/A.

Usage Guidelines

This command creates an MLAG peer switch association structure.

You must use a unique name for the peer switch. If you attempt to create an MLAG peer with a name that already exists, the following error message is displayed:

ERROR: MLAG peer with specified name already exists

Example

The following command creates a peer switch structure switch101:

```
# create mlag peer switch101
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

create mpls rsvp-te path

```
create mpls rsvp-te path path_name
```

Description

Creates an RSVP-TE routed path resource.

Syntax Description

<i>path_name</i>	Identifies the path within the switch. The character string must begin with an alphabetic character, and may contain up to 31 additional alphanumeric characters.
------------------	---

Default

N/A.

Usage Guidelines

This command creates an RSVP-TE path resource.

The `path_name` parameter must begin with an alphabetical character and may contain alphanumeric characters and underscores (`_`), but it cannot contain spaces. The maximum allowed length for a name is 32 characters. For name creation guidelines and a list of reserved names, see *Object Names* in the [Switch Engine 32.3 User Guide](#).

The maximum number of configurable paths is 255.



Note

The RSVP-TE LSP is not signaled along the path until an LSP is created and then configured with the specified `path_name`.

Example

The following example creates an RSVP-TE path:

```
create mpls rsvp-te path path598
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

create mpls rsvp-te lsp

```
create mpls rsvp-te lsp lsp_name destination ipaddress
```

Description

Creates internal resources for an RSVP-TE LSP.

Syntax Description

<i>lsp_name</i>	Specifies a name for the LSP you are creating. The character string must begin with an alphabetic character, and may contain up to 31 additional alphanumeric characters.
<i>ipaddress</i>	Specifies the endpoint of the LSP.

Default

N/A.

Usage Guidelines

This command creates internal resources for an RSVP-TE LSP.

The LSP name must begin with an alphabetical character and may contain alphanumeric characters and underscores (`_`), but it cannot contain spaces. The maximum allowed length for a name is 32 characters.

For name creation guidelines and a list of reserved names, see *Object Names* in the [Switch Engine 32.3 User Guide](#).

The ipaddress specifies the endpoint of the LSP. The LSP is not signaled until a path is specified for the LSP using the `configure mpls rsvp-te lsp lsp_name add path` command. When multiple LSPs are configured to the same destination, IP traffic is load-shared across active LSPs that have IP transport enabled. The maximum number of RSVP-TE LSPs that can be created is 1024.



Note
The LSP must be created before it can be configured.

Example

The following command creates an RSVP-TE LSP:

```
create mpls rsvp-te lsp lsp598 destination 11.100.100.8
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support *MPLS* as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

create mpls rsvp-te profile fast-reroute

```
create mpls rsvp-te profile profile_name fast-reroute
```

Description

Creates an LSP container to hold FRR configuration parameters.

Syntax Description

<i>profile_name</i>	Specifies a name for the new RSVP-TE fast-reroute profile. The character string must begin with an alphabetic character and may contain up to 31 additional alphanumeric characters.
---------------------	--

Default

N/A.

Usage Guidelines

A profile name must begin with an alphabetical character and may contain alphanumeric characters and underscores (_), but it cannot contain spaces. The maximum allowed length for a name is 32 characters. For name creation guidelines and a list of reserved names, see *Object Names* in the [Switch Engine 32.3 User Guide](#).

Example

The following command creates a new FRR profile named frrprofile:

```
create mpls rsvp-te profile frrprofile fast-reroute
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available only on the platforms that support *MPLS* as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

create mpls rsvp-te profile

```
create mpls rsvp-te profile profile_name {standard}
```

Description

Creates configured RSVP-TE profile with the specified profile name.

Syntax Description

<i>profile_name</i>	Identifies the RSVP-TE profile. The character string must begin with an alphabetic character and may contain up to 31 additional alphanumeric characters.
standard	The standard option differentiates this command version from the command that creates a fast-reroute profile. If you do not specify an option, a standard RSVP-TE profile is created.

Default

N/A.

Usage Guidelines

This command creates 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.

A profile name must begin with an alphabetical character and may contain alphanumeric characters and underscores (_), but it cannot contain spaces. The maximum allowed length for a name is 32 characters. For name creation guidelines and a list of reserved names, see *Object Names* in the [Switch Engine 32.3 User Guide](#).

Example

The following command creates an RSVP-TE profile:

```
create mpls rsvp-te profile prof598
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support *MPLS* as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

create mpls static lsp

```
create mpls static lsp lsp_name destination ipaddress
```

Description

Creates internal resources for a static LSP and assigns a name to the LSP.

Syntax Description

<i>lsp_name</i>	Identifies the LSP to be created.
<i>ipaddress</i>	Specifies the endpoint of the LSP.

Default

N/A.

Usage Guidelines

An LSP name must begin with an alphabetical character and may contain alphanumeric characters and underscores (_), but it cannot contain spaces. The maximum allowed length for a name is 32 characters. For name creation guidelines and a list of reserved names, see *Object Names* in the [Switch Engine 32.3 User Guide](#).

Example

The following command creates a static LSP:

```
create mpls static lsp lsp598 destination 11.100.100.8
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

create msdp mesh-group

```
create msdp mesh-group mesh-group-name {vr vrname}
```

Description

Creates an [MSDP](#) mesh-group.

Syntax Description

<i>mesh-group-name</i>	Specifies the name for the MSDP mesh-group.
<i>vrname</i>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

N/A.

Usage Guidelines

A mesh-group is a group of MSDP peers with fully meshed MSDP connectivity. Create a mesh-group to:

- Reduce SA message flooding.
- Simplify peer-RPF flooding.

SA messages received from a peer in a mesh-group are not forwarded to other peers in the same mesh-group, which reduces SA message flooding.

A mesh group name must begin with an alphabetical character and may contain alphanumeric characters and underscores (_), but it cannot contain spaces. The maximum allowed length for a name is 32 characters. For name creation guidelines and a list of reserved names, see *Object Names* in the [Switch Engine 32.3 User Guide](#).

Example

The following example creates a mesh-group called "verizon":

```
create msdp mesh-group verizon
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

create msdp peer

```
create msdp peer remoteaddr {remote-as remote-AS} {vr vrname}
```

Description

Creates an *MSDP* peer.

Syntax Description

<i>remoteaddr</i>	Specifies the IP address of the MSDP router to configure as an MSDP peer.
<i>remote-AS</i>	Specifies the autonomous system (AS) number of the MSDP peer. This optional parameter is deprecated in ExtremeXOS 12.1, though the option is still available in the CLI for backward compatibility. The software ignores this parameter.
<i>vrname</i>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

N/A.

Usage Guidelines

The *BGP* route database is used by MSDP to determine the AS number for the peer. You can display the AS number (which can be a 2-byte for 4-byte AS number) using the command:

```
show msdp [peer {detail} | {peer} remoteaddr] {vrvrname}.
```

Example

The following example creates an MSDP peer:

```
create msdp peer 192.168.45.43 remote-as 65001
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

create netlogin local-user

```
create netlogin local-user user-name {encrypted} encrypted_password |
  password } {vlan-vsa [[{tagged | untagged} [vlan_name] | vlan_tag]]}
  {security-profile security_profile}
```

Description

Creates a local network login user name and password.

Syntax Description

<i>user-name</i>	Specifies a new local network login user name. User names must have a minimum of 1 character and a maximum of 32 characters.
encrypted	The encrypted option is used by the switch to encrypt the password. Do not use this option through the command line interface (CLI).
<i>password</i>	Specifies a local network login user password. Passwords must have a minimum of 0 characters and a maximum of 32 characters.
tagged	Specifies that the client be added as tagged.
untagged	Specifies that the client be added as untagged.
<i>vlan_name</i>	Specifies the name of the destination <u>VLAN</u> .
<i>vlan_tag</i>	Specifies the VLAN ID, tag, of the destination VLAN.
<i>security_profile</i>	Specifies a security profile string during account creation.

Default

N/A.

Usage Guidelines

Use this command to create a local network login account and to configure the switch to use its local database for network login authentication. This method of authentication is useful in the following situations:

- If both the primary and secondary (if configured) RADIUS servers timeout or are unable to respond to authentication requests.
- If no RADIUS servers are configured.
- If the RADIUS server used for network login authentication is disabled.

If any of the above conditions are met, the switch checks for a local user account and attempts to authenticate against that local account.

Extreme Networks recommends creating a maximum of 64 local accounts. If you need more than 64 local accounts, we recommend using RADIUS for authentication. For more information about RADIUS authentication, see the [Switch Engine 32.3 User Guide](#).

You can also specify the destination VLAN to enter upon a successful authentication.

**Note**

If you do not specify a password or the keyword `encrypted`, you are prompted for one.

Additional Requirements

This command applies only to the web-based and MAC-based modes of network login. 802.1X network login does not support local database authentication.

You must have administrator privileges to use this command. If you do not have administrator privileges, the switch displays a message similar to the following:

```
This user does not have permissions for this command.
```

User names are not case-sensitive. Passwords are case-sensitive. User names must have a minimum of 1 character and a maximum of 32 characters. Passwords must have a minimum of 0 characters and a maximum of 32 characters. If you use RADIUS for authentication, we recommend that you use the same user name and password for both local authentication and RADIUS authentication.

If you attempt to create a user name with more than 32 characters, the switch displays the following messages:

```
%% Invalid name detected at '^' marker. %% Name cannot exceed 32 characters.
```

If you attempt to create a password with more than 32 characters, the switch displays the following message after you re-enter the password:

```
Password cannot exceed 32 characters
```

Modifying an Existing Account

To modify an existing local network login account, use the following command:

```
configure netlogin local-user user-name {vlan-vsa [{tagged | untagged}  
[vlan_name | vlan_tagw]] | none}}
```

Displaying Local Network Login Accounts

To display a list of local network login accounts on the switch, including VLAN information, use the following command:

```
show netlogin local-users
```

Example

The following command creates a local network login user name and password:

```
create netlogin local-user megtest
```

After you enter the local network login user name, press [Enter]. The switch prompts you to enter a password (the switch does not display the password):

```
password:
```

After you enter the password, press [Enter]. The switch then prompts you to re-enter the password:

```
Reenter password:
```

The following command creates a local network login user name, password, and associates a destination VLAN with this account:

```
create netlogin local-user accounting vlan-vsa blue
```

As previously described, the switch prompts you to enter and confirm the password.

History

This command was first available in ExtremeXOS 11.2.

The **vlan-vsa** parameter and associated options were added in ExtremeXOS 11.3.

The **security-profile** parameter was added in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

create ntp key

```
create ntp key keyid [md5 | sha256] {encrypted encrypted_key_string |  
  key_string}
```

Description

Enables an NTP key for an NTP session.

Syntax Description

<i>keyid</i>	Specifies the key ID as a value from 1 to 65534.
<i>key_string</i>	Specifies an alphanumeric key string, from 5 to 20 numbers or characters, or a combination of both.
md5	Specifies MD5 authentication type. This authentication type is not allowed when FIPS mode is on.
sha256	Specifies SHA-265 authentication type.

Default

N/A.

Usage Guidelines

N/A.

Example

The following command enables an NTP key using RSA Data Security, Inc. [MD5](#) Message-Digest Algorithm encryption on the switch:

```
create ntp key 1 md5 oklahoma
```

History

This command was first available in ExtremeXOS 12.7.

Key string length was changed to 20 in ExtremeXOS 30.3.

SHA-256 option was added in ExtremeXOS 30.4.

The **md5** option is not allowed when FIPS mode is on in ExtremeXOS 30.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

create ospf area

```
create ospf area area-identifier
```

Description

Creates an [OSPF](#) area.

Syntax Description

<i>area-identifier</i>	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 first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [***Switch Engine 32.3 Feature License Requirements***](#) document.

create ospfv3 area

```
create ospfv3 area area_identifier
```

Description

Creates an OSPFv3 area.

Syntax Description

area_identifier	Specifies an OSPFv3 area, a four-byte, dotted decimal number.
-----------------	---

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 a non-backbone OSPFv3 area:

```
create ospfv3 area 1.2.3.4
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

create policy access-list

```
create policy access-list list_dot_rule {matches [ {app-signature
  group group name name} {ether ether {mask ether_mask}} {icmp6type
  icmp6type {mask icmp6_mask}} {icmptype icmptype {mask icmp_mask}}
  {ipdestsocket ipdestsocket {mask ipdest_mask}} {ipfrag} {ipproto
  ipproto {mask ipproto_mask}} {ipsourcesocket ipsourcesocket {mask
  ipsrc_mask}} {iptos iptos {mask iptos_mask}} {ipttl ipttl {mask
  ipttl_mask}} {tcpdestportIP tcpdestportIP {mask tcpdest_mask}}
  {tcpsourceportIP tcpsourceportIP {mask tcpsrc_mask}} {udpdestportIP
  udpdestportIP {mask udpdest_mask}} {udpsourceportIP udpsourceportIP
  {mask udpsrc_mask}} ] } {actions [ {cos cos} {drop | forward}
  {mirror-destination control_index} {syslog}]}
```

Description

Creates policy access-list match criteria.

Syntax Description

access-list	Specifies access-list rule model to select multiple match criteria per rule.
<i>list_dot_rule</i>	Specifies the access-list name and rule name in the format <i>list_name.rule_name</i> .
matches	Selects up to 5 match criteria.
app-signature	Associates an application signature to a policy profile. Note: Not supported on the ExtremeSwitching X435 series switches.
group	Associates an application signature group to a policy profile.
<i>group</i>	Defines the application signature group name.
name	Associates an application signature name to a policy profile.
<i>name</i>	Defines the name assigned to the application signature (range 1–32).
ether	Selects the type field in Ethernet II packet.
<i>ether</i>	Defines the type field in Ethernet II packet (data: 0–65535 or 0x0–0xFFFF; mask: 1–16).
mask	Selects a mask.
<i>ether_mask</i>	Selects the number of most significant bits to match data value (range 1–16).
icmp6type	Selects ICMPv6 type.code.

<i>icmp6type</i>	Defines the ICMPv6 type.code (data: 123.456 (dotted-decimal) or AB-CD (dashed-hexadecimal)).
<i>icmp6_mask</i>	Specifies the number of most significant bits to match data value (range 1-16).
icmptype	Selects an ICMP type.code.
<i>icmptype</i>	Specifies the ICMP type.code - (data: a.b; mask: 1-16).
<i>icmp_mask</i>	Specifies the number of most significant bits to match data value (range 1-16).
ipdestsocket	Specifies a destination IP address with optional post-fixed port or port-range.
<i>ipdestsocket</i>	Defines the destination IP address with optional post-fixed port or port-range - (data: a.b.c.d [:ab (0-65535) [-cd (0-65535)]]; mask: 1-48,64).
<i>ipdest_mask</i>	Specifies the number of most significant bits to match data value (range 1-64).
ipfrag	Selects IP fragmentation flag.
ipproto	Specifies protocol field in IP packet.
<i>ipproto</i>	Defines the protocol field in IP packet (data: 0-255 or 0x0-0xFF; mask: 1-8). IPv4 only (ICMP).
<i>ipproto_mask</i>	Specifies the number of most significant bits to match the data value (range 1-8).
ipsourcesocket	Specifies the source IP address with optional post-fixed port or port-range.
<i>ipsourcesocket</i>	Defines the source IP address with optional post-fixed port or port-range - (data: a.b.c.d [:ab (0-65535) [-cd (0-65535)]]; mask: 1-48, 64).
<i>ipsrc_mask</i>	Specifies the number of most significant bits to match data value (range 1-64).
iptos	Specifies IPv4 type of service/IPv6 traffic class field.
<i>iptos</i>	Defines the IPv4 type of service/IPv6 traffic class field (data: 0-255; mask: 1-8).
<i>iptos_mask</i>	Specifies the number of most significant bits to match data value (range 1-8).
ipttl	Specifies IP time to live.
<i>ipttl</i>	Defines the IP time to live (data: 0-255 or 0x0-0xFF; mask:1-8).
<i>ipttl_mask</i>	Specifies the number of most significant bits to match data value (range 1-8).
tcpdestportIP	Specifies TCP port/port-range destination with optional post-fix IPv4 address.
<i>tcpdestportIP</i>	Defines the TCP port/port-range destination with optional post-fix IPv4 address (data: ab [-cd] [:c.d.e.f]); mask: 1-64).
<i>tcpdest_mask</i>	Specifies the number of most significant bits to match data value (range 1-64).

tcpsourceportIP	Specifies TCP port/port-range source with optional post-fix IPv4 address.
<i>tcpsourceportIP</i>	Defines the TCP port/port-range source with optional post-fix IPv4 address (data: ab [-cd] [:c.d.e.f]); mask: 1-64).
<i>tcpsrc_mask</i>	Specifies the number of most significant bits to match data value (range 1-64).
udpdestportIP	Specifies UDP port/port-range destination with optional post-fix IPv4 address.
<i>udpdestportIP</i>	Defines the UDP port/port-range destination with optional post-fix IPv4 address (data: ab [-cd] [:c.d.e.f]); mask:1-64).
<i>udpdest_mask</i>	Specifies the number of most significant bits to match data value (range 1-64).
udpsourceportIP	Specifies UDP port/port-range source with optional post-fix IPv4 address.
<i>udpsourceportIP</i>	Defines the UDP port/port-range source with optional post-fix IPv4 address (data: ab [-cd] [:c.d.e.f]).
<i>udpsrc_mask</i>	Specifies the number of most significant bits to match data value (range 1-64).
actions	Specifies selecting one or more actions to occur when there is a match.
cos	Specifies Class of Service (CoS) as an action.
<i>cos</i>	Defines the CoS (0-255), or -1 for no CoS, or CoS with no forwarding behavior to remove the existing forwarding settings.
drop	Specifies dropping any packets that match this rule.
forward	Specifies forwarding any packets that match this rule.
mirror-destination	Specifies mirroring any packets that match this rule.
<i>control_index</i>	Defines which mirror destination control index (1-4).
syslog	Enables, disables, or prohibits Syslog using event Policy.LogRuleHit on first rule use.

Default

N/A.

Usage Guidelines

To use this command, the policy rule model must be set to access-list (use command `configure policy rule-model [access-list | hierarchical]`).

The following combinations are not allowed:

- ipfrag with icmp, tcp, udp or ip with port rules
- tcp/udp source rules with ipSrc rule with port
- tcp/udp rules dest rule with ipDest rule with port
- icmp with tcp, udp or ip with port rules

Example

The following example creates the policy access list "ACL1.ace3" with match criteria of IP source address "10.1.1.1" and mask "32" with the action to forward with Class of Service level "2":

```
# create policy access-list ACL1.ace3 matches ipsource 10.1.1.1 mask 32 actions forward
cos 2
```

History

This command was first available in ExtremeXOS 30.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

create policy access-list action-set

```
create policy access-list action-set set-id [{drop | forward} {cos cos}
{mirror-destination control_index} {syslog}]
```

Description

Creates a pre-defined set for use in RADIUS Change of Authorization (CoA) and Extreme Dynamic ACL via Radius VSA 232.

Syntax Description

access-list	Specifies access-list features.
action-set	Defines a set of actions that can be applied to multiple sets of match conditions.
<i>set-id</i>	Identifies the global action-set ID (range 1-63).
drop	Specifies dropping any packets that match this rule.
forward	Specifies forwarding any packets that match this rule.
cos	Specifies setting Class of Service (CoS).
<i>cos</i>	Specifies the CoS value: 0-255, or -1 for no CoS, or CoS with no forwarding behavior to remove forwarding behavior.
mirror-destination	Specifies setting a mirror destination control index.
<i>control_index</i>	Specifies setting the mirror destination control index (1-4).
syslog	Specifies Syslog logging using event Policy.LogRuleHit when first rule use occurs.

Default

N/A.

Usage Guidelines

You can view your configurations made with this command using the `show policy access-list action-set {set_id}` command.

Example

The following example creates an action set "1" with CoS level of 3 and Syslog behavior:

```
# create policy access-list action-set 1 cos 3 syslog
```

This command will be accepted only if the mode is set to access-list and slices must be shared by entering the command **configure policy slices shared 2** and **configure policy slices tci-overwrite 2**.

History

This command was first available in ExtremeXOS 30.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

create ports group

```
create ports group port_group
```

Description

This command creates a generic port-group name that can be associated with a list of ports. The *port_group* option could be implemented in configure or show commands that currently accept a *port_list*. The *QoS* commands are expanded to accept the *port_group* option. QoS commands that use port groups are updated automatically if the ports group is removed or if ports are added or removed from the group.

Syntax Description

<i>port_group</i>	Specifies a port group name.
-------------------	------------------------------

Default

N/A.

Usage Guidelines

Use this command to create a generic port-group name to be associated with a list of ports.

Example

```
create ports group testGroup
```

History

This command was first available in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

create private-vlan

```
create private-vlan name {vr vr_name}
```

Description

Creates a PVLAN framework with the specified name.

Syntax Description

<i>name</i>	Specifies a name for the new PVLAN.
<i>vr_name</i>	Specifies the VR in which the PVLAN is created.

Default

N/A.

Usage Guidelines

The PVLAN is a framework that links network and subscriber VLANs; it is not an actual VLAN.

A private VLAN name must begin with an alphabetical character and may contain alphanumeric characters and underscores (_), but it cannot contain spaces. The maximum allowed length for a name is 32 characters. For private VLAN naming guidelines and a list of reserved names, see *Object Names* in the [Switch Engine 32.3 User Guide](#).

If no VR is specified, the PVLAN is created in the default VR context.

Example

The following example creates a PVLAN named "companyx":

```
create private-vlan companyx
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms that support the Private VLAN feature. For features and the platforms that support them, see the [Switch Engine 32.3 Feature License Requirements](#) document.

create process executable

```
create process name executable exe {start [auto | on-demand]} {node
  node} {vr vr-name} {description description} {arg1 {arg2 { arg3
  { arg4 { arg5 { arg6 { arg7 { arg8 { arg9 }}}}}}}}}}
```

Description

Adds a C executable process compiled using the C-based SDK.

Syntax Description

process	User application process.
<i>name</i>	Name of the user application process. Range 1-31.
executable	Executable.
<i>exe</i>	Name of the executable relative to /usr/local/cfg.
start	Startup behavior.
auto	Create a daemon process and start it immediately
on-demand	Create a run-to-completion process and use \"start process\".
node	Node in stack in which to create the process.
<i>node</i>	Primary node, backup node, or both (default is primary).
vr	Virtual router in which to start the process.
<i>vr-name</i>	Virtual router name (Default is <u>VR-Mgmt</u>).
<i>description</i>	Description.
<i>arg1-9</i>	Variable value.

Default

VR-Mgmt is the default VR used if not specified.

If no selection is made, the process runs on-demand.

If no node is selected, the default is the primary node.

Usage Guidelines

The executable must be uploaded to `/usr/local/cfg` using the normal mechanisms (for example, TFTP).

Fields are provided by the user and passed directly into an `epmrc` entry. Not all `epmrc` fields are available.

This command adds C executable processes. To add a Python module, use the [create process python-module](#) on page 1661 command.

A process must first exist on the primary node if you are creating it only on the backup node. If a process already exists on the primary node, you cannot create it on both the primary and secondary node. Also, if the backup node is down, a new process cannot be created on it.

Example

```
create process foo_userd executable foobard start auto vr VR-Default description "Run
foobard on the default VR" "arg1" "arg2 with spaces"
```

The following error is displayed if an attempt is made to create a process with an invalid name:

```
Error: Process name %s is invalid. Process names must begin with a letter, contain only
alphanumeric and
"_" characters, and be less than 32 characters long.
```

History

This command was first available in ExtremeXOS 15.7.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

create process python-module

```
create process name python-module python-module {start [auto | on-
demand]} {node node} {vr vr-name} {description description} {arg1
arg2 {arg3 {arg4 {arg5 {arg6 {arg7 {arg8 {arg9}}}}}}}}}
```

Description

Adds a Python module process.

Syntax Description

process	User application process.
<i>name</i>	Name of the user application process. Range 1-31.
python-module	The Python module to import and run.
<i>python-module</i>	The module relative to <code>/usr/local/cfg</code> .

start	Startup behavior.
auto	Create a daemon process and start it immediately.
on-demand	Create a run-to-completion process and use <code>\start process\</code> .
node	Node in stack in which to create the process.
<i>node</i>	Primary node, backup node, or both (default is primary).
vr	Virtual router in which to start the process.
<i>vr-name</i>	Virtual router name.
<i>description</i>	Description.
<i>arg1-9</i>	Variable value.

Default

VR-Mgmt is the default VR used if not specified.

If no selection is made, the process runs on-demand.

If no node is selected, the default is the primary node.

Usage Guidelines

The executable must be uploaded to `/usr/local/cfg` using the normal mechanisms (for example, TFTP).

From EPM's perspective, a Python-based process is an instance of the "expy" executable with some arguments, namely the Python module.

This command adds a Python module. To add a C executable processes, use the [create process executable](#) on page 1660 command.

A process must first exist on the primary node if you are creating it only on the backup node. If a process already exists on the primary node, you cannot create it on both the primary and secondary node. Also, if the backup node is down, a new process cannot be created on it.

Example

The following are examples of create process python-module commands.

```
python-module foo_program start auto vr vr-default
create process foo_user1 python-module "foo.run" "arg1 to foo.main"
create process foo_user2 python-module "foo.noargs.needed"
create process foo_user3 python-module "foo.daemon" start auto "arg1 to foo.main"
```

The following error is displayed if an attempt is made to create a process with an invalid name:

```
Error: Process name %s is invalid. Process names must begin with a letter, contain only
alphanumeric and
"_" characters, and be less than 32 characters long.
```

History

This command was first available in ExtremeXOS 15.7.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

create protocol

```
create protocol {filter} filter_name
```

Description

Creates a user-defined protocol filter.

Syntax Description

filter	Specifies a protocol filter.
<i>filter_name</i>	Specifies a protocol filter name. The protocol filter name can have a maximum of 31 characters.

Default

N/A.

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 configure protocol command. To assign it to a VLAN, use the `configure {vlan} vlan_name protocol {filter} filter_name` command.

Example

The following command creates a protocol named "my_filter", and a protocol filter named "my_other_filter":

```
create protocol "my_filter"  
create protocol filter "my_other_filter"
```

History

This command was first available in ExtremeXOS 10.1.

The **filter** keyword was added in ExtremeXOS 15.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

create qosprofile

```
create qosprofile [QP2 | QP3 | QP4 | QP5 | QP6 | QP7]
```

Description

Creates a QoS profile.

Syntax Description

QP2 . . . QP7	Specifies the QoS profile you want to create.
---------------	---

Default

N/A.

Usage Guidelines

ExtremeSwitching series switches allow dynamic creation and deletion of QoS profiles QP2 to QP7. Creating a QoS profile dynamically does not cause loss of traffic.

QoS profiles QP1 and QP8 are part of the default configuration and cannot be deleted. You must create a QoS profile in the range of QP2 to QP7 before you can configure it or assign it to traffic groups.

Qos profile QP7 cannot be created in a SummitStack; this queue is reserved for control traffic.



Note

The sFlow application uses QP2 to sample traffic on SummitStack and ExtremeSwitching series switches; any traffic grouping using QP2 can encounter unexpected results when sFlow is enabled on these specific devices.

Example

The following command creates QoS profile QP3:

```
create qosprofile qp3
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

create snmp trap

```
create snmp trap severity severity event EventName msg
```

Description

Creates and sends an *SNMP* trap containing the information defined in the command.

Syntax Description

<i>severity</i>	Specifies one of the eight severity levels defined in the ExtremeXOS software. Enter one of the following values: critical, error, warning, notice, info, debug-summary, debug-verbose, debug-data.
<i>EventName</i>	Specifies the event name. Enter a name using alphanumeric characters.
<i>msg</i>	Specifies a message. Enter the message using alphanumeric characters.

Default

N/A.

Usage Guidelines

None.

Example

The following example sends a trap of severity info for event AAA with the message user XYZ logged in:

```
create snmp trap severity info event AAA "user XYZ logged in"
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

create sshd2 key-file

```
create sshd2 key-file {host-key | user-key} key_name
```

Description

Creates a file for the user-key or host-key.

Syntax Description

host-key	Specifies the name of the host-key.
user-key	Specifies the name of the user-key.
<i>key_name</i>	Specifies the name of the public key.

Default

N/A.

Usage Guidelines

This command is used to write the user or the host public key in a file. The key files will be created with a .ssh file extension; this enables the administrator to copy the public key files to another server.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

create sshd2 user-key

```
create sshd2 user-key key_name key {subject subject} {comment comment}
```

Description

Creates a user key.

Syntax Description

<i>key_name</i>	Specifies the name of the public key.
<i>key</i>	Specifies the key. Note: The key cannot have any spaces in it.
<i>subject</i>	Specifies the subject.
<i>comment</i>	Specifies the comment (an optional field).

Default

N/A.

Usage Guidelines

This command is used to enter, or cut and paste, your public key. You can also enter the public key into the switch by using the SCP or SFTP client that is connected to the switch.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

create stpd

```
create stpd stpd_name {description stpd-description}
```

Description

Creates a user-defined STPD.

Syntax Description

<i>stpd_name</i>	Specifies a user-defined STPD name to be created. May be up to 32 characters in length.
<i>stpd-description</i>	Specifies an <u>STP</u> domain description string.

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.
- Maximum BPDU age—20 seconds.
- Hello time—2 seconds.
- Forward delay—15 seconds.
- Operational mode—802.1D.
- Rapid Root Failover—disabled.
- Default Binding Mode (encapsulation mode)—Ports in the default STPD (s0) are in 802.1d mode. Ports in user-created STPDs are in emistp mode.
- Maximum hop count (when configured for MSTP)—20 hops.
- STP domain description string—empty.

Usage Guidelines

The maximum length for a name is 32 characters. Names can contain alphanumeric characters and underscores (_) but cannot be any reserved keywords, for example, stp or stpd. Names must start with an alphabetical character, for example, a, Z. For name creation guidelines and a list of reserved names, see [Object Names](#) on page 11.

Each STPD name must be unique and cannot duplicate any other named STPDs on the switch. If you are uncertain about the STPD names on the switch, use the `show stpd` command to view the STPD names.

You can, however, re-use names across multiple categories of switch configuration. For example, you can use the name Test for an STPD and a VLAN. If you use the same name, we recommend that you specify the appropriate keyword when configuring the STPD. If you do not specify the appropriate keyword, the switch displays a message similar to the following:

```
%% Ambiguous command: "configure Test"
```

To view the names of the STPDs on the switch, enter configure and press [Tab]. Scroll to the end of the output to view the names.

The maximum length for an STPD description is 180 characters. The description must be in quotes if the string contains any spaces. To display the description, use the `show stpd stpd_name` command.

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 ExtremeXOS 10.1.

The STPD description option was added in ExtremeXOS 12.4.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

create tunnel 6to4

```
create tunnel tunnel_name 6to4 source source-address
```

Description

Creates an IPv6-to-IPv4 (6to4) tunnel.

Syntax Description

<i>tunnel_name</i>	Specifies an IPv6 tunnel.
<i>source-address</i>	Specifies an IPv4 address for the tunnel.

Default

N/A.

Usage Guidelines

This command will create a new IPv6-to-IPv4 (also known as a 6to4 tunnel), and add it to the system. Only one 6to4 tunnel can be configured on any particular VR.

The tunnel name must be unique and cannot overlap the same name space as VLANs, other tunnels, or VRs. The name must begin with an alphabetical character and may contain alphanumeric characters and underscores (_), but it cannot contain spaces. The maximum allowed length for a name is 32 characters. For name creation guidelines and a list of reserved names, see *Object Names* in the [Switch Engine 32.3 User Guide](#).

The source address of the tunnel must be one of the IPv4 addresses already configured on the switch. You cannot remove an IPv4 address from the switch if a tunnel that uses it still exists.

Example

The following example creates the 6to4 tunnel "link35" with source address 192.168.10.1:

```
create tunnel link35 6to4 source 192.168.10.1
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 interworking feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

create tunnel gre destination source

```
create tunnel tunnel_name gre destination destination-address source  
       source-address {vr vr_name} {payload-vr payload_vr_name}
```

Description

Allows you to add a GRE tunnel. This command is in-line with adding an ipv6-in-ipv4 tunnel.

Syntax Description

gre	Generic Routing Encapsulation tunnel.
<i>destination-address</i>	IPv4 destination address of the tunnel.
<i>source-address</i>	IPv4 source address of the tunnel.
vr	Specifies configuring the tunnel on specific VR.
<i>vr_name</i>	Specifies which VR to configure the tunnel on. This VR is the VR of the tunnel itself, where <i>source-address</i> is configured. It is the underlay/delivery. This is the source outer IP address. The default is the VR of the current context.
payload-vr	Specifies a tunnel payload VR (VR of the tunnel interface).
<i>payload_vr_name</i>	Specifies the tunnel payload VR name. It is the VR of the L3 interface of the tunnel, that is, the payload/overlay. The default is the VR of the tunnel.

Default

No GRE tunnels exist in the system.

When adding tunnels, the VR of the current context is the default unless otherwise specified.

By default, the payload VR is the VR of the tunnel.

Usage Guidelines

Use this command to add a GRE tunnel.

Example

```
create tunnel myGREtunnel gre destination 10.0.0.2 source 10.0.0.1
```

History

This command was first available in ExtremeXOS 15.3.

Ability to configure GRE tunnels on user VRs was added in ExtremeXOS 31.3.

Platform Availability

This command is available on the platforms listed for the GRE feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

create tunnel ipv6-in-ipv4

```
create tunnel tunnel_name ipv6-in-ipv4 destination destination-address
source source-address
```

Description

Creates an IPv6-in-IPv4 (6in4) tunnel.

Syntax Description

<i>tunnel_name</i>	Specifies an IPv6 tunnel.
<i>source-address</i>	Specifies an IPv4 address for the tunnel.

Default

N/A.

Usage Guidelines

This command creates a new IPv6-in-IPv4 (otherwise known as a configured tunnel or a 6in4 tunnel) and add it to the system. A maximum of 255 tunnels (including one 6to4 tunnel) can be configured on the system.

The tunnel name must be unique and cannot overlap the same name space as VLANs, other tunnels, or VRs. The name must begin with an alphabetical character and may contain alphanumeric characters and underscores (_), but it cannot contain spaces. The maximum allowed length for a name is 32 characters. For name creation guidelines and a list of reserved names, see *Object Names* in the [Switch Engine 32.3 User Guide](#).

The source address of the tunnel must be one of the IPv4 addresses already configured on the switch. You cannot remove an IPv4 address from the switch if a tunnel is still exists that uses it.

Example

The following example creates the 6in4 tunnel "link39" with destination address 10.10.10.10 and source address 192.168.10.15:

```
create tunnel link39 ipv6-in-ipv4 destination 10.10.10.10 source 192.168.10.15
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 interworking feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

create upm profile

```
create upm profile profile-name
```

Description

Creates a new profile of a specified type.

Syntax Description

<i>profile-name</i>	Specifies the UPM profile to be created.
---------------------	--

Default

N/A.

Usage Guidelines

Use this command to create a profile and name it. The maximum profile size is 5000 characters.

A UPM profile name must begin with an alphabetical character and may contain alphanumeric characters and underscores (_), but it cannot contain spaces. The maximum allowed length for a name is 32 characters. For name creation guidelines and a list of reserved names, see [Object Names](#) on page 11.

There is a limited capability to edit the profile with this command. If you enter a period (.) as the first and the only character on a line, you terminate the editing of the file. Use the command: `edit upm profile profile-name` for block mode capability.

Example

The following example shows how to create a profile named "P2":

```
# create upm profile p2
enable port 2:*
disable port 3:1
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the Universal Port feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

create upm timer

```
create upm timer timer-name
```

Description

Creates and names a UPM timer.

Syntax Description

<i>timer-name</i>	Specifies the name of the UPM timer to be created.
-------------------	--

Default

N/A.

Usage Guidelines

You can create UPM timers with a name. A profile can be associated with eight timers, but a timer can be bound to only one profile at any point in time. You can create a maximum of 32 timers. A name space for the timers is available to help when you are typing the commands.

A UPM timer name must begin with an alphabetical character and may contain alphanumeric characters and underscores (_), but it cannot contain spaces. The maximum allowed length for a name is 32 characters. For name creation guidelines and a list of reserved names, see *Object Names* in the [Switch Engine 32.3 User Guide](#).

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the Universal Port feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

create virtual-network

```
create virtual-network vn_name {flooding [standard | explicit-remotes |  
    multicast {group grpipaddress}}
```

Description

This command creates a virtual network instance in ExtremeXOS. The virtual network instance maps to a bridge instance within ExtremeXOS.

Syntax Description

<i>vn_name</i>	Alphanumeric string identifying the Virtual Network to be created.
flooding	Configure flooding method for unknown-destination frames.

standard	Standard L2 flooding behavior to remote endpoints and tenant ports.
explicit-remotes	Explicitly configured flooding to remote endpoints with standard L2 flooding to tenant ports.
multicast	Multicast flooding to remote endpoints with standard L2 flooding to tenant ports.
group	Configure multicast group for flooding of unknown-destination frames.
<i>grpipaddress</i>	IPv4 multicast group address to be used for flooding (Automatically assigned if unspecified).

Default

Standard.

Usage Guidelines

For auto-assigning multicast groups, you must configure the following command: **configure virtual-network multicast group**.

This bridge instance is not dependent on the overlay encapsulation scheme. The virtual network name can be a maximum of 32 characters. The current restrictions on naming objects in ExtremeXOS apply. Virtual network names are added to a new namespace within ExtremeXOS. Virtual networks may use one of two flooding methods for flooding to remote endpoints. The “standard” mode offers handling of unknown destination frames very similar to standard Layer 2. The unknown destination frames are flooded to all local ports and remote endpoints. The “explicit-remotes” mode offers granular control of which remote endpoints receive certain types of unknown destination frames. Different remote endpoint sets may be configured for; broadcast, unknown unicast, and unknown multicast. These sets are configured with `create fdb` and `configure fdb` commands

Example

The following example creates the virtual network "my_virtual_network":

```
create virtual-network my_virtual_network
```

The following example deletes the virtual network "my_virtual_network":

```
delete virtual-network my_virtual_network
```

History

This command was first available in ExtremeXOS 21.1.

Multicast flooding support was made available in ExtremeXOS 31.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

create virtual-network remote-endpoint vxlan ipaddress

```
create virtual-network remote-endpoint vxlan ipaddress ipaddress {vr
    vr_name}
```

Description

This command creates a remote endpoint.

Syntax Description

<i>ipaddress</i>	A remote endpoint IP address.
vr	VR/VRF instance the remote endpoint is associated with
<i>vr_name</i>	An existing VR/VRF name.

Default

VR-Default

Usage Guidelines

This command is useful when you want to explicitly add a remote endpoint in addition to the ones learnt dynamically (OSPF extensions). In flood mode explicit, you must create a remote-endpoint using this command, if the configurations on remote-endpoint (like monitor) need to be saved to the configuration. Otherwise, the configuration will be lost after the switch reboots.

Example

To create a remote endpoint:

```
create virtual-network vxlan remote-endpoint ipaddress 1.2.3.4
```

History

This command was first available in ExtremeXOS 21.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

create virtual-router

To create virtual routers, use the following command:

```
create virtual-router vr-name {type [vrf | vpn-vrf {vr parent_vr_name}]}
```

To create local-only virtual routers (ExtremeSwitching 5320-24T/P and 5320-16P only), use the following command:

```
create virtual-router vr-name local-only {type [vrf | vpn-vrf {vr
  parent_vr_name}}]}
```

Description

Use the create virtual-router command to create a user VR or VRF.

Syntax Description

<i>vr-name</i>	Virtual router name.
type	Specifies the type of virtual router you are creating.
local-only	Specifies local-only VR. For ExtremeSwitching 5320-24T/P and 5320-16P series switches and stacks only.
vrf	Specifies that you are creating a new L3 or IP routing domain.
vpn-vrf	Specifies that you are creating a new L3 or IP routing domain that supports L3VPNs. Not supported on Universal platforms.
<i>parent_vr_name</i>	Specifies the parent VR that supports the VRF you are creating.

Default

If no **type** is specified, then the default is to create a user virtual router. A virtual router creates separate L3 Routing Domains.

If *parent_vr_name* parameter is not specified, the VRF will be created under the VR of the current CLI context. The default is VR-Default.

Usage Guidelines

All VRFs are created under default VR or a user created VR. VPN-VRFs can be created in any VR but for L3VPNs to work, VPN-VRFs should be created under a parent VR where MPLS is configured. There is a single namespace maintained by the configuration manager and it contains VRs and VRFs. Hence the name for a VR or a VRF must be unique in ExtremeXOS.

A VR or VRF name must begin with an alphabetical character and may contain alphanumeric characters and underscores (_), but it cannot contain spaces. The maximum allowed length for a name is 31 characters. The name must be unique among the object names on the switch, and the name is case insensitive. For information on VR and VRF name guidelines and a list of reserved names, see *Object Names* in the [Switch Engine 32.3 User Guide](#).

When a new VR is created, by default, no ports are assigned, no VLAN interface is created, and no support for any routing protocols is added. A protocol process is started in the parent VR when a protocol instance is added to a VRF. If you do not specify a VR type, this command creates a user VR.

VRFs are supported as children of user VRs or VR-Default. If a *parent_vr_name* is specified when a VRF is created, the new VRF is created under that parent, provided that the parent supports VRFs. If

no parent is specified, the VRF is assigned to the VR for the current VR context, or to VR-Default if the current VR context does not support VRFs.

**Note**

To support Layer 3 VPNs, a VPN VRF must be created under the VR that supports MPLS. The software supports MPLS on only one VR.

Starting with ExtremeXOS 22.6, you can create "local-only" virtual routers that have separate logical IP lookup tables used only for IP packets to or from the switch's local IP addresses. This feature is only applicable for ExtremeSwitching X440-G2 and X620 series switches and stacks with these switches. All other platforms support separate logical IP lookup tables in hardware, so "local-only" is not specified.

Example

The following example creates the VR "vr-acme":

```
create virtual-router vr-acme
```

The following example creates the non-VPN VRF vrf1:

```
create virtual-router vrf1 type vrf
```

The following example creates the local-only VR "vrl" (on 5320-24T/P and 5320-16P series switches only):

```
create virtual-router vrl local-only
```

History

This command was first available in ExtremeXOS 11.0.

Support for non-VPN VRFs was added in ExtremeXOS 12.5.

Support for VPN VRFs was added in ExtremeXOS 12.6.0-BGP.

Support for L3 VPN VRFs was added in ExtremeXOS 15.3.

Support for local-only VRs was added in ExtremeXOS 22.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

create vlan

```
create vlan [ vlan_name {tag tag} | vlan_list ] {description vlan-  
description } {vr name }
```

Description

Creates a named VLAN.

Syntax Description

<i>vlan_name</i>	Specifies a VLAN name (up to 32 characters).
<i>vlan_list</i>	Specifies a VLAN list of IDs.
tag	Specifies a value to use as an 802.1Q tag. The valid range is from 2 to 4095.
<i>vlan-description</i>	Specifies a VLAN description (up to 64 characters) that appears in <code>show vlan</code> commands and can be read from the ifAlias MIB object for the VLAN.
<i>name</i>	<p>Specifies a VR or virtual routing and forwarding (VRF) instance in which to create the VLAN.</p> <p>Note: User-created VRs are supported only on the platforms listed for this feature in the Switch Engine 32.3 Feature License Requirements document. On switches that do not support user-created VRs, all VLANs are created in <u>VR-Default</u> and cannot be moved.</p>

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.

A VLAN named Mgmt exists on switches that have management ports:

- It initially contains the management port(s) the switch.
- It is assigned the next available internal VLANid as an 802.1Q tag.

If you do not specify the VR, the VLAN is created in the current VR.

If the VLAN description contains one or more space characters, you must enclose the complete name in double quotation marks.

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 `configure vlan` commands to configure the VLAN to your needs.

Internal VLANids are assigned automatically using the next available VLANid starting from the high end (4094) of the range.

The VLAN name can include up to 32 characters. VLAN names must begin with an alphabetical letter, and only alphanumeric, underscore (_), and hyphen (-) characters are allowed in the remainder

of the name. VLAN names cannot match reserved keywords. For more information on VLAN name requirements and a list of reserved keywords, see *Object Names* in the [Switch Engine 32.3 User Guide](#).

**Note**

If you use the same name across categories (for example, [STPD](#) and EAPS names), we recommend that you specify the identifying keyword as well as the actual name. If you do not use the keyword, the system may return an error message.

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

You must use mutually exclusive names for:

- VLANs
- VMANs
- Ipv6 tunnels
- BVLANS
- SVLANS
- CVLANS

**Note**

The VLAN description is stored in the ifAlias MIB object.

If you do not specify a VR when you create a VLAN, the system creates that VLAN in the default VR (VR-Default). The management VLAN is always in the management VR ([VR-Mgmt](#)).

Once you create VRs, ExtremeXOS allows you to designate one of these as the domain in which all your subsequent configuration commands, including VLAN commands, are applied. If you create VRs, ensure that you are creating the VLANs in the desired virtual-router domain.

**Note**

User-created VRs are supported only on the platforms listed for this feature in the [Switch Engine 32.3 Feature License Requirements](#) document.. On switches that do not support user-created VRs, all VLANs are created in VR-Default and cannot be moved.

Example

The following example creates a VLAN named accounting on the current VR:

```
create vlan accounting description "Accounting Dept"
```

History

This command was first available in ExtremeXOS 10.1.

The **vr** option was added in ExtremeXOS 11.0.

The *vlan-description* option was added in ExtremeXOS 12.4.4.

The *vlan_list* variable was added in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

create vm image

```
create vm vm_name image image_file {memory memory_size} {cpus num_cpus}
    {slot slot_ID} {vnc [none | vnc_display]}
```

Description

Creates a guest virtual machine (VM) from a disk image file.

Syntax Description

vm	Designates creating a virtual machine.
<i>vm_name</i>	Specifies the VM name.
image	Designates using a disk image file to create the VM.
<i>image_file</i>	Specifies the disk image file to use in qcow2 or any QEMU-supported (including VMDK) format.
memory	Designates specifying the amount of RAM allocated to the VM.
<i>memory_size</i>	Specifies the amount of RAM (in MB) allocated to the VM. The default is 4,096.
cpus	Designates specifying the number of CPUs to allocate to the VM.
<i>num_cpus</i>	Specifies the number of CPUs to allocate to the VM. Range is 1–2. The default is 1.
slot	Specifies assigning the VM to run on a slot.
<i>slot_ID</i>	Specifies the slot number that the VM will run on.
vnc	Specifies providing a display number for VNC access.
none	Disables VNC access (default).
<i>vnc_display</i>	Specifies the VNC screen number. Range is 0–15.

Default

The default memory size to run the VM on is 4,096 MB.

The default number of CPUs to allocate to the VM is one.

By default, VNC access is disabled.

Usage Guidelines

The disk image must be a qcow2 or any QEMU-compatible file.

If the VM storage device has not been initialized when this command is run, you are prompted to run the `clear vm storage` command to initiate partitioning, file system creation, and initialization of the file/directory structure on the device.

The Integrated Application Hosting (IAH) feature requires the Solid State Storage Device SSD-120.

For the VNC display number (or screen number), you can use the values from 0 to 15. These correspond to TCP ports 5,900 to 5,915. Multiple VMs can be configured with the same VNC display, but VMs configured with the same display number cannot run at the same time. A VM cannot be started if the VNC port is already in use. For security reasons, the VNC display is only accessible using SSH tunnel. VNC cannot be configured on non-primary stack nodes.

Example

The following example creates a VM called "vm1" with disk image file "my_file" with 2,000 MB as the amount of RAM allocated to the VM:

```
# create vm vm1 image my_file memory 2000
```

History

- This command was first available in ExtremeXOS 30.3.
- VMDK format support was added in ExtremeXOS 30.4.
- VNC capability and support for any QEMU-compatible disk was added in ExtremeXOS 30.5.
- Stacking support was added in ExtremeXOS 30.6.

Platform Availability

This command is available on all platforms that support the Integrated Application Hosting (IAH) feature and have a Core or Premier license installed. For a list of platforms that support the IAH feature and for information about licenses, see the [Switch Engine 32.3 Feature License Requirements](#).

create vm ova

```
create vm vm_name ova ova_file {memory memory_size} {cpus num_cpus}
      {slot slot_ID} {vnc [none | vnc_display]}
```

Description

Creates a guest virtual machine (VM) from an Open Virtual Appliance (OVA) file.

Syntax Description

vm	Designates creating a virtual machine.
<i>vm_name</i>	Specifies the VM name.
ova	Designates using an OVA file to create the VM.

<i>ova_file</i>	Specifies the OVA file to use.
memory	Designates specifying the amount of RAM allocated to the VM.
<i>memory_size</i>	Specifies the amount of RAM (in MB) allocated to the VM. The default is 4,096.
cpus	Designates specifying the number of CPUs to allocate to the VM.
<i>num_cpus</i>	Specifies the number of CPUs to allocate to the VM. Range is 1–2. The default is 1.
slot	Specifies assigning the VM to run on a slot.
<i>slot_ID</i>	Specifies the slot number that the VM will run on.
vnc	Specifies providing a display number for VNC access.
none	Disables VNC access (default).
<i>vnc_display</i>	Specifies the VNC screen number. Range is 0–15.

Default

The default memory size to run the VM on is 4,096 MB.

The default number of CPUs to allocate to the VM is one.

By default, VNC access is disabled.

Usage Guidelines

If the VM storage device has not been initialized when this command is run, you are prompted to run the `clear vm storage` command to initiate partitioning, file system creation, and initialization of the file/directory structure on the device.

Compatibility issues may occur when using third-party OVA files. The image format qcow2 is generally more reliable.

The Integrated Application Hosting (IAH) feature requires the Solid State Storage Device SSD-120.

For the VNC display number (or screen number), you can use the values from 0 to 15. These correspond to TCP ports 5,900 to 5,915. Multiple VMs can be configured with the same VNC display, but VMs configured with the same display number cannot run at the same time. A VM cannot be started if the VNC port is already in use. For security reasons, the VNC display is only accessible using SSH tunnel. VNC cannot be configured on non-primary stack nodes.

Example

The following example creates a VM called "vm1" with OVA file "my_ova" with 2,000 MB as the amount of RAM allocated to the VM:

```
# create vm vm1 ova my_ova memory 2000
```

History

This command was first available in ExtremeXOS 30.3.

VNC capability was added in ExtremeXOS 30.5.

Stacking support was added in ExtremeXOS 30.6.

Platform Availability

This command is available on all platforms that support the Integrated Application Hosting (IAH) feature and have a Core or Premier license installed. For a list of platforms that support the IAH feature and for information about licenses, see the [Switch Engine 32.3 Feature License Requirements](#).

create vman

```
create vman [vman-name | vman_list] {learning-domain} {vr vr_name}
```

Description

Creates a VMAN.

Syntax Description

<i>vman-name</i>	Specifies a VMAN name using up to 32 characters.
<i>vman_list</i>	Specifies the VMAN tag range or VMAN Tag List (Ex: 2-4 or 2,3).
learning-domain	Specifies that this VMAN is a learning domain, which supports inter-VMAN forwarding.
vr	Specifies a virtual router.
<i>vr_name</i>	Specifies a virtual router name. Note: User-created VRs are supported only on the platforms listed for this feature in the Switch Engine 32.3 Feature License Requirements document. On switches that do not support user-created VRs, all VLANs are created in <u>VR-Default</u> and cannot be moved.

Default

N/A.

Usage Guidelines

For information on VMAN name requirements and a list of reserved keywords, see [Object Names](#) on page 11. You must use mutually exclusive names for:

- VLANs
- VMANs
- IPv6 tunnels

The keyword `learning-domain` enables you to create a VMAN that serves as a learning domain for inter-VMAN forwarding.

If you do not specify the virtual router, the VMAN is created in the current virtual router. After you create the VMAN, you must configure the VMAN tag and add the ports that you want.

Example

The following example creates a VMAN named "fred":

```
create vman fred
```

History

This command was first available in ExtremeXOS 11.0.

The `vman_list` variable was added in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

create vm-tracking local-vm

```
create vm-tracking local-vm mac-address mac {name name | ipaddress
ipaddress vpp vpp_name | vlan-tag tag {vr vr_name}}
```

Description

Creates a local VM database entry to be used for VM MAC local authentication, with optional parameters.

Syntax Description

<i>mac</i>	Specifies the MAC address for the VM. This must match the MAC address configured on the VM and be unique among the locally configure VM addresses.
<i>name</i>	Specifies a name to represent this VM in show vm-tracking command display.
<i>ipaddress</i>	Specifies the IP address for the VM. This must match the IP address configured on the VM.
<i>vpp_name</i>	Specifies the virtual port profile to apply for the local VM.
<i>tag</i>	<u>VLAN</u> tag between 1 and 4094.
<i>vr_name</i>	Virtual router name.

Default

N/A.

Usage Guidelines

A VM name can include up to 32 characters. VM names must begin with an alphabetical letter, and only alphanumeric, underscore (_), and hyphen (-) characters are allowed in the remainder of the name. VM names cannot match reserved keywords. For more information on VM name requirements and a list of reserved keywords, see [Object Names](#).

The following command creates a VM entry named VM1 in the local VM database:

```
# create vm-tracking local-vm mac-address 00:E0:2B:12:34:56 name VM1
```

The following command creates a VM entry and assigns IP address 10.10.2.2 to the entry:

```
# create vm-tracking local-vm mac-address 00:E0:2B:12:34:57 ip-address 10.10.2.2
```

The following command creates a VM entry and assigns VPP vpp1 to it:

```
# create vm-tracking local-vm mac-address 00:E0:2B:12:34:58 vpp vpp1
```

History

This command was first available in ExtremeXOS 12.5.

The ingress-vpp and egress-vpp options were replaced with the vpp option in ExtremeXOS 12.6.

The vlan-tag and vr-name options were added in ExtremeXOS 15.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

create vm-tracking vpp

```
create vm-tracking vpp vpp_name
```

Description

Creates a Local VPP (LVPP).

Syntax Description

<i>vpp_name</i>	Specifies a name for the new VPP.
-----------------	-----------------------------------

Default

N/A.

Usage Guidelines

A VPP name can include up to 32 characters. VPP names must begin with an alphabetical letter, and only alphanumeric, underscore (_), and hyphen (-) characters are allowed in the remainder of the name.

VPP names cannot match reserved keywords. For more information on VPP name requirements and a list of reserved keywords, see [Object Names](#) on page 11.

Example

The following example creates a VPP named vpp1:

```
# create vm-tracking vpp vpp1
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

create vpls fec-id-type pseudo-wire

create vpls *vpls_name* **fec-id-type** **pseudo-wire** *pwid*



Note

This command has been replaced with the following command: `create l2vpn [vpls vpls_name | vpwsvpws vpws_name] fec-id-type pseudo-wire pwid` . This command is still supported for backward compatibility, but it will be removed from a future release, so we recommend that you start using the new command.

Description

Creates a VPLS instance with the specified *vpls_name*.

Syntax Description

<i>vpls_name</i>	Identifies the VPLS within the switch (character string). The <i>vpls_name</i> string must begin with an alphabetic character, and may contain up to 31 additional alphanumeric characters.
<i>pwid</i>	Specifies a PW ID. Must be a non-zero 32-bit value that has network-wide significance.

Default

For the VPLS dot1q tag, the default value is exclude.

Usage Guidelines

This command creates a VPLS instance with the specified *vpls_name*. Each VPLS represents a separate virtual switch instance (VSI).

The `vp1s_name` parameter must begin with an alphabetical character and may contain alphanumeric characters and underscores (`_`), but it cannot contain spaces. The maximum allowed length for a name is 32 characters. For name creation guidelines and a list of reserved names, see *Object Names* in the [Switch Engine 32.3 User Guide](#).

Each VPLS is a member of a single VPN and each VPN may have only one associated VPLS per switch. External to the switch, each VPN has an identifier.

Any non-zero 32-bit value that has network-wide significance can be specified for the identifier. This `pwid` is used on all pseudowires in the VPLS.



Note
The switch's LSR ID must be configured before a VPLS can be created.

Example

This example creates a VPLS with 99 as the PW ID:

```
create vpls vpls1 fec-id-type pseudo-wire 99
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

create vrrp group

```
create vrrp group group_name
```

Description

This command defines a [VRRP](#) group to operate in high-scale mode.

Syntax Description

group	Specifies setting up a VRRP group for high-scale mode.
<i>group_name</i>	Specifies the VRRP group name.

Default

None.

Example

The following example creates a VRRP group called "vrrp1".

```
create vrrp group vrrp1
```

History

This command was first available in ExtremeXOS 22.2.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the VRRP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

create vrrp vlan vrid

```
create vrrp vlan [vlan_name | vlan_list] vrid [vridval | vrid_list]
```

Description

Creates a VRRP instance on the switch.

Syntax Description

<i>vlan_name</i>	Specifies the name of a VRRP <u>VLAN</u> .
<i>vlan_list</i>	VLAN list (1-4,094).
<i>vridval</i>	Specifies a VRID for the VRRP instance. The value can be in the range of 1-255.
<i>vrid_list</i>	List of virtual router IDs (1-255).

Default

N/A.

Usage Guidelines

VRRP Router IDs can be used across multiple VLANs. You can create multiple VRRP routers on different VLANs. VRRP router IDs need not be unique to a specific VLAN.



Note
The total number of supported VRRP router instances is dependent on the switch hardware. For more information, see the ExtremeXOS Release Notes.

Before configuring any VRRP router parameters, you must first create the VRRP instance on the switch. If you define VRRP parameters before creating the VRRP, you might see an error similar to the following:

```
Error: VRRP VR for vlan vrrp1, vrid 1 does not exist.
Please create the VRRP VR before assigning parameters.
Configuration failed on backup MSM, command execution aborted!
```

If this happens, create the VRRP instance and then configure its parameters.

Example

The following command creates a VRRP router on VLAN vrrp-1, with a VRRP router ID of 1:

```
create vrrp vlan vrrp-1 vrid 1
```

History

This command was first available in ExtremeXOS 10.1.

VLAN and VR list options added in ExtremeXOS 22.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the VRRP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

create xml-notification target url

```
create xml-notification target new-target url url {vr vr_name} {user
  [none | user]} {encrypted-auth encrypted-auth} {queue-size queue-
  size}
```

Description

Creates the Web server target in the XML client.

Syntax Description

<i>new-target</i>	Specifies a name for the target being created.
<i>url</i>	Specifies the Web server URL.
<i>vr_name</i>	Specifies the name of the virtual router over which the XML client process can connect to the Web server.
<i>user</i>	Specifies the name of the user.

<i>encrypted-auth</i>	Specifies the encrypted user authentication string.
<i>queue-size</i>	Specifies, in numeric format, the size of the buffer that stores incoming events from ExtremeXOS.

Default

N/A.

Usage Guidelines

Use this command to create the Web server target in the XML client process.



Note

You cannot enter a password in the CLI directly. It is a two-step process similar to creating a user account in ExtremeXOS.

Example

The following command creates a target target2 on http://10.255.129.22:8080/xos/webservice with a queue size of 100:

```
create xml-notification target target2 url http://10.255.129.22:8080/xos/webservice queue-size 100
```

History

This command was first available in ExtremeXOS 12.4.

The virtual router option was added in ExtremeXOS 12.4.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

delete access-list

```
delete access-list dynamic_rule
```

Description

Deletes a dynamic *ACL*.

Syntax Description

<i>dynamic_rule</i>	Specifies the dynamic ACL name.
---------------------	---------------------------------

Default

N/A.

Usage Guidelines

This command deletes a dynamic ACL rule. Before you delete a dynamic ACL, it must be removed from any interfaces it is applied to. Use the [configure access-list delete](#) command to remove the ACL from an interface.

Example

The following command deletes the dynamic ACL icmp-echo:

```
delete access-list icmp-echo
```

History

This command was first available in ExtremeXOS 11.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

delete access-list network-zone

```
delete access-list network-zone zone_name
```

Description

This command is used to delete a network-zone and all configurations that belong to that zone.

Syntax Description

<i>zone_name</i>	Network-zone name
------------------	-------------------

Default

N/A.

Usage Guidelines

Use this command to delete a network-zone and all configurations belonging to that zone.

Example

```
Switch# delete access-list network-zone zone1
```

If the user tries to delete a network-zone that is bound with one or more policy files, the following error message will be displayed, and the command will be rejected.

```
Switch # delete access-list network-zone zone1
Error: Network Zone "zone1" - Unable to delete zone. Zone has one
or more policies.
```

History

This command was first available in ExtremeXOS 15.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

delete access-list zone

delete access-list zone *name*

Description

Deletes an *ACL* zone.

Syntax Description

<i>name</i>	Specifies the zone name.
-------------	--------------------------

Default

N/A.

Usage Guidelines

This command deletes an ACL zone. You must remove all applications from a zone before you can delete the zone. To delete an application from a zone, use the command [configure access-list zone *name* delete application *appl-name*](#).

You cannot delete the default zones.

Example

The following command deletes the zone my_zone:

```
delete access-list zone my_zone
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

delete account

```
delete account name
```

Description

Deletes a specified user account.

Syntax Description

<i>name</i>	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. The system must have one administrator account; the command will fail if an attempt is made to delete the last administrator account on the system.

To ensure security, change the password on the default account, but do not delete it. The changed password will remain intact through configuration uploads and downloads.

If you must delete the default account, first create another administrator-level account.

Example

The following command deletes account John2:

```
delete account John2
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

delete auto-peering

```
delete auto-peering
```

Description

This command deletes auto-peering (either OSPF or BGP), removing all of the auto-peering configuration. This command deletes the VLAN list, loopback, and BGP configuration created with enabling auto-peering.

Syntax Description

This command has no arguments or variables.

Default

N/A

Usage Guidelines



Important

Deleting auto-peering when executed on a large leaf-spine topology causes massive change in the network with many route withdrawals and updates.

Example

The following example deletes auto-peering:

```
# delete auto-peering
```

History

This command was first available in ExtremeXOS 22.5.

Ability to delete OSPFv2 Auto-peering was added in ExtremeXOS 30.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

This feature requires the Advanced Edge license. For more information about licenses, see the [Switch Engine 32.3 Feature License Requirements](#).

delete bgp evpn instance

```
delete bgp evpn instance evpn_instance_name
```

Description

Deletes an EVPN instance.

Syntax Description

bgp	BGP capability.
evpn	EVPN protocol.
instance	Specifies deleting an EVPN instance.
<i>evpn_instance_name</i>	Name of the EVPN instance.

Default

N/A.

Example

The following example deletes an EVPN instance named "my_evpn":

```
# delete bgp evpn instance my_evpn
```

History

This command was first available in ExtremeXOS 30.2.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

delete bgp neighbor

```
delete bgp neighbor [remoteaddr | all]
```

Description

Deletes one or all *BGP* neighbors.

Syntax Description

<i>remoteaddr</i>	Specifies the IPv4 or IPv6 address of the BGP neighbor to be deleted.
all	Specifies all IPv4 and IPv6 neighbors.

Default

N/A.

Usage Guidelines

You can use global unicast remote addresses to delete all BGP peer types. You can use link-local remote address to delete only EBGp single-hop peers.

Example

The following command deletes the specified IPv4 BGP neighbor:

```
delete bgp neighbor 192.168.1.17
```

The following command deletes the specified IPv6 BGP neighbor:

```
delete bgp neighbor fe80::204:96ff:fe1e:a8f1%vlan1
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

delete bgp peer-group

```
delete bgp peer-group peer-group-name
```

Description

Deletes a peer group.

Syntax Description

<i>peer-group-name</i>	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 outer:

```
delete bgp peer-group outer
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

delete cfm domain

```
delete cfm domain domain
```

Description

Deletes the specified maintenance domain (MD) from the switch, as well as all configuration setting related to this MD.

Syntax Description

<i>domain</i>	Enter the name of the domain you want to delete.
---------------	--

Default

N/A.

Usage Guidelines

This command deletes all configuration settings related to the domain—for example, all MAs, MIPs, and MEPs—as well as the domain itself.

Example

The following command deletes the domain atlanta (as well as all settings related to this domain):

```
delete cfm domain atlanta
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

delete cfm segment

```
delete cfm segment [segment_name | all]
```

Description

Deletes one or all CFM segments.

Syntax Description

<i>segment_name</i>	An alpha-numeric string identifying the segment name.
all	Specifies all CFM segments.

Default

N/A.

Usage Guidelines

Use this command to delete one or all CFM segments.

Example

The following example deletes the CFM segment "segment-new":

```
delete cfm segment segment-new
```

History

This command was first available in ExtremeXOS 12.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

delete database

```
delete database database_name
```

Description

Deletes an Automation Edge remote VXLAN network identifier (VNI)-device database.

Syntax Description

database	Deletes a remote VNI-database.
<i>database_name</i>	Specifies the name of the database to delete.

Default

N/A.

Usage Guidelines

N/A.

Example

The following example deletes a database called "database1":

```
# delete database database1
```

History

This command was first available in ExtremeXOS 31.1 as a demonstration feature.

This command is fully supported in ExtremeXOS 31.2.

Platform Availability

This command is available on ExtremeSwitching 5420 and 5520 series switches.

delete eaps shared-port

```
delete eaps shared-port ports
```

Description

Deletes an *EAPS* shared port on a switch.

Syntax Description

<i>ports</i>	Specifies the port number of the Common Link port.
--------------	--

Default

N/A.

Usage Guidelines

None.

Example

The following command deletes shared port 1:1.

```
delete eaps shared-port 1:1
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms with the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and what licenses are appropriate for this feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

delete eaps

```
delete eaps name
```

Description

Deletes the *EAPS* domain with the specified name.

Syntax Description

<i>name</i>	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 ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

delete erps

delete erps *ring-name*

Description

Deletes an *ERPS* ring.

Syntax Description

<i>ring-name</i>	Alphanumeric string that identifies the ERPS ring.
------------------	--

Default

N/A.

Usage Guidelines

Use this command to delete an ERPS ring.

Example

The following command deletes an ERPS ring named “ring1”:

```
delete erps ring1
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

delete esrp

```
delete esrp esrpDomain
```

Description

Deletes the *ESRP* domain with the specified name.

Syntax Description

<i>esrpDomain</i>	Specifies the name of an ESRP domain to be deleted.
-------------------	---

Default

N/A.

Usage Guidelines

You must first disable an ESRP domain before you delete it. To disable an ESRP domain, use the `disable esrp` command.

You do not have to remove the master or member VLANs from an ESRP domain before you delete it. When you delete an ESRP domain, All VLANs are automatically removed from the domain.

For ESRP domains configured of type VPLS-redundancy, you need to unconfigure all associated VPLS instances from the ESRP domain using the `unconfigure vpls redundancy` command before deleting the domain.

Example

The following command deletes ESRP domain esrp1 from the switch:

```
delete esrp esrp1
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

delete fdb mac-tracking entry

```
delete fdb mac-tracking entry [mac_addr | all]
```

Description

Deletes a MAC address from the MAC address tracking table.

Syntax Description

<i>mac_addr</i>	Specifies a device MAC address, using colon-separated bytes.
all	Specifies that all MAC addresses are to be deleted from the MAC address tracking table.

Default

The MAC address tracking table is empty.

Usage Guidelines

None.

Example

The following example deletes a MAC address from the MAC address tracking table:

```
delete fdb mac-tracking entry 00:E0:2B:12:34:56
```

History

This command was first available in ExtremeXOS 12.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

delete fdb

```
delete fdb [all | mac_address [vlan vlan_name ] | vxlan { vr vr_name }  
            {ipaddress} remote_ipaddress ] | broadcast vlan vlan_name vxlan  
            { vr vr_name } {ipaddress} remote_ipaddress | unknown-multicast  
vlan vlan_name vxlan { vr vr_name } {ipaddress} remote_ipaddress  
            | unknown-unicast vlan vlan_name vxlan { vr vr_name } {ipaddress}  
            remote_ipaddress ]
```

Description

Deletes one or all permanent *FDB* entries.

Syntax Description

all	Specifies all FDB entries.
<i>mac_address</i>	Specifies a device MAC address, using colon-separated bytes.
<i>vlan_name</i>	Specifies the specific <u>VLAN</u> name.
broadcast	Forwarding destination(s) for broadcast traffic.
unknown-unicast	Forwarding destination(s) for unknown unicast traffic.
<i>unknown-unicast</i>	Forwarding destination(s) for unknown multicast traffic.
vxlan	The MAC address is reachable through a VXLAN Tunnel.
vr	VR/VRF instance the IPv4 address is configured on.
<i>vr_name</i>	An existing VR/VRF name.
ipaddress	Configure the IP address of the remote tunnel endpoint to which the MAC needs to be bound.
<i>remote_ipaddress</i>	IPv4 address of the remote tunnel endpoint.

Default

N/A.

Usage Guidelines

In ExtremeXOS 21.1, this command was extended to delete a remote VTEP as a destination to a MAC address. Three new tokens “broadcast”, “unknown-multicast” and “unknown-unicast” have been added to this command. When you want to specify a destination to forward all broadcast or unknown unicast traffic on that VLAN, these token are used. For “broadcast”, “unknown-multicast” and “unknown-unicast” only remote VTEPs (and not port_list or blackhole) can be specified in this release of ExtremeXOS. These entries can only be created when the virtual-network is in explicit-remote flooding mode.

Example

The following example deletes a permanent entry from the FDB:

```
delete fdb 00:E0:2B:12:34:56 vlan marketing
```

The following example deletes all permanent entries from the FDB:

```
delete fdb all
```

History

This command was first available in ExtremeXOS 11.0.

In ExtremeXOS 12.3, the **fdb** keyword was introduced as an alias to the **fdbentry** keyword to avoid interference with the syntax of the MAC-Tracking feature commands. Both keywords execute; however, the syntax helper (tab completion) does not recognize the **fdbentry** keyword.

Three new tokens “broadcast”, “unknown-multicast” and “unknown-unicast” were added to this command in ExtremeXOS 21.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

delete flow-redirect

```
delete flow-redirect flow_redirect_name
```

Description

Deletes the named flow redirection policy.

Syntax Description

<i>flow_redirect_name</i>	Specifies the name of the flow redirection policy.
---------------------------	--

Default

N/A.

Usage Guidelines

Use this command to delete a named flow-redirection policy. Before it can be deleted, all nexthop information must be deleted, otherwise an error message is displayed.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

delete flowmon collector

```
delete flowmon collector collector_name
```

Description

Deletes a Flow Monitor collector.

Syntax Description

collector	Specifies to send flow information to a collector.
<i>collector_name</i>	Specifies the name of the created collector. Range is 32 characters.

Default

N/A.

Usage Guidelines

The system will reject any attempt to delete a collector that does not exist.

Example

The following command deletes a collector with the name 'src-ipv4-address':

```
# delete flowmon collector src-ipv4-address
```

History

This command was first available in ExtremeXOS 32.2.

Platform Availability

This command is available on ExtremeSwitching 5520 and 5420 series switches.

delete flowmon group

```
delete flowmon group group_name
```

Description

Deletes a Flow Monitor group.

Syntax Description

group	Specifies the Flow Monitor group.
<i>group_name</i>	Specifies the assigned name of the Flow Monitor group. Range is 32 characters.

Default

N/A.

Usage Guidelines

An enabled group can't be deleted and the system will reject any attempt to do so. Groups must be disabled before deleting.

The system will also reject any attempt to delete a group that does not already exist.

Example

The following command deletes a group with the name 'max-flow-age':

```
# delete flowmon group max-flow-age
```

History

This command was first available in ExtremeXOS 32.2.

Platform Availability

This command is available on ExtremeSwitching 5520 and 5420 series switches.

delete flowmon key

```
delete flowmon key key_name
```

Description

Deletes a Flow Monitor key.

Syntax Description

key	Specifies the Flow Monitor key.
<i>key_name</i>	Specifies the assigned name of the Flow Monitor key. Range is 32 characters.

Default

N/A.

Usage Guidelines

If a key has been added to a group and the delete command is used, the key is automatically deleted from the group. The related template key is disassociated and only deleted when there are no more keys or groups referencing it.

The system will reject any attempt to delete a key that does not already exist.

Example

The following command deletes a key with the name 'src-ipv4-addr':

```
# delete flowmon key src-ipv4-addr
```

History

This command was first available in ExtremeXOS 32.2.

Platform Availability

This command is available on ExtremeSwitching 5520 and 5420 series switches.

delete identity-management role

```
delete identity-management role {role_name | all}
```

Description

Deletes one or all roles.

Syntax Description

<i>role_name</i>	Specifies a name of an existing role to delete.
all	Specifies that all roles are to be deleted.

Default

N/A.

Usage Guidelines

Any policy applied to users of a deleted role gets reverted. The users are placed under one of the other roles based on their attributes. Parent and child relationships to other roles are also deleted. For example, all child roles under the deleted role become orphans and hence they and their descendants no longer inherit the policies of the deleted role.

Example

The following example deletes the role named India-Engr:

```
* Switch.99 # delete identity-management role "India-Engr"
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

delete ip nat rule

```
delete ip nat rule rule_name
```

Description

Deletes an IP Network Address Translation (NAT) rule.

Syntax Description

ip	Specifies Internet Protocol (IP).
nat	Specifies NAT.
rule	Specifies deleting a NAT rule.
<i>rule_name</i>	Specifies the NAT rule name to delete.

Default

N/A.

Usage Guidelines

To create a rule, run the command `create ip nat rule rule_name type [source-nat | napt | destination-napt]`.

Example

The following example deletes the NAT rule named "rule1"

```
# delete ip nat rule rule1
```

History

This command was first available in ExtremeXOS 31.2.

Platform Availability

This command is available on the switches that support the IP NAT feature. For information about which switches support this and other features, see the [Switch Engine 32.3 Feature License Requirements](#) document.

delete isis area

```
delete isis area [all | area_name]
```

Description

This command disables and deletes the specified IS-IS router process in the current virtual router.

Syntax Description

all	Deletes all IS-IS router processes.
<i>area_name</i>	Specifies the name of the IS-IS router process to be deleted.

Default

None.

Usage Guidelines

All configuration for the specified router is lost. All routes learned from this router process are purged from the routing tables.

Example

The following command deletes the IS-IS process named areax:

```
delete isis area areax
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Premier license as described in the [**Switch Engine 32.3 Feature License Requirements**](#) document.

delete keychain

```
delete keychain keychain_name
```

Description

This command deletes a keychain.

Syntax Description

<i>keychain_name</i>	Defines a name for the keychain. The range is 1-31.
----------------------	---

Default

N/A.

Usage Guidelines

Use this command to delete a key from the keychain.

Example

The following command deletes a keychain:

```
delete keychain ospfv3-keys
```

History

This command was first available in ExtremeXOS 31.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

delete l2pt profile

```
delete l2pt profile profile_name
```

Description

Deletes an L2PT profile.

Syntax Description

l2pt	Deletes a Layer 2 protocol tunneling profile.
profile	Profile that defines L2PT configuration for L2 protocols.
<i>profile_name</i>	Specifies a profile name (maximum 32 characters).

Default

Disabled.

Usage Guidelines

Use this command to delete an L2PT profile.

Example

The following example deletes *my_l2pt_prof* that is currently in use by a service:

```
delete l2pt profile my_l2pt_prof
```

The following example deletes *my_l2pt_prof* that is not associated with any service:

```
delete l2pt profile my_l2pt_prof
```

History

This command was first available in ExtremeXOS 15.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

delete l2vpn

```
delete l2vpn [vppls [vppls_name | all] | vpws [vpws_name | all]]
```

Description

Deletes the specified VPLS or VPWS.

Syntax Description

<i>vppls_name</i>	Identifies the VPLS within the switch (character string).
<i>vpws_name</i>	Identifies the VPWS within the switch (character string).
all	Specifies all VPLS or VPWS instances.

Default

N/A.

Usage Guidelines

All PWs established to VPLS or VPWS peers are terminated.

The **l2vpn** keyword is introduced in ExtremeXOS Release 12.4 and is required when deleting a VPWS. For backward compatibility, the **l2vpn** keyword is optional when deleting a VPLS. However, this keyword will be required in a future release, so we recommend that you use this keyword for new configurations and scripts.

Example

This commands deletes the VPLS myvpls:

```
delete vpls myvpls
```

This commands deletes the VPWS myvpws:

```
delete l2vpn vpws myvpws
```

History

This command was first available in ExtremeXOS 11.6.

The l2vpn and vpws keywords were first available in ExtremeXOS 12.4.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

delete ldap domain

```
delete ldap domain [domain_name | all]
```

Description

This command is used to delete one or all LDAP domains.

When an LDAP domain is deleted, all LDAP servers added under that domain are also deleted. Also all LDAP configurations done for that domain are deleted.

Syntax Description

<i>domain_name</i>	Name of the LDAP domain that wil be deleted.
--------------------	--

Default

N/A.

Usage Guidelines

Use this command to delete one or all LDAP domains.

When an LDAP domain is deleted, all LDAP servers added under that domain are also deleted. All LDAP configurations for that domain are also deleted.

Example

This command deletes the LDAP domain sales.XYZCorp.com

```
delete ldap domain sales.XYZCorp.com
```

History

This command was first available in ExtremeXOS 15.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

delete log filter

```
delete log filter [filter-name | all]
```

Deletes a log filter with the specified name.

Syntax Description

<i>filter-name</i>	Specifies the filter to delete.
all	Specifies that all filters, except DefaultFilter, are to be deleted

Default

N/A.

Usage Guidelines

This command deletes the specified filter, or all filters except for the filter DefaultFilter. The specified filter must not be associated with a target. To remove that association, associate the target with DefaultFilter instead of the filter to be deleted, using the following command:

```
configure log target target filter DefaultFilter
```

Example

The following command deletes the filter named fdb2:

```
delete log filter fdb2
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

delete log target upm

```
delete log target upm {upm_profile_name}
```

Description

Deletes the specified UPM log target.

Syntax Description

<i>upm_profile_name</i>	Specifies the name of the UPM log target to be deleted.
-------------------------	---

Default

N/A.

Usage Guidelines

This command deletes the log target and any configurations applied to that target. To disable a target and retain the target configuration, use the following command:

```
disable log target upm {upm_profile_name}.
```

Example

The following command deletes the UPM log target testprofile1:

```
delete log target upm testprofile1
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the Universal Port feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

delete log target xml-notification

```
delete log target xml-notification xml_target_name
```

Description

Deletes a Web server target.

Syntax Description

<i>xml_target_name</i>	Specifies the name of the xml notification target.
------------------------	--

Default

N/A.

Usage Guidelines

Use this command to delete a Web server target.

Example

The following command deleted the Web server target target2:

```
delete log target xml-notification target2
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

delete macsec connectivity-association

```
delete macsec connectivity-association ca_name
```

Description

Deletes a previously created connectivity-association (CA) object that holds MAC Security (MACsec) key authentication data.

Syntax Description

connectivity-association	Secures connectivity provided between MACsec stations.
<i>ca_name</i>	Selects the CA to delete.

Default

N/A.

Usage Guidelines

Prior to deletion, ports assigned to the CA must be removed with the `configure macsec connectivity-association` *ca_name* [**pre-shared-key** {*ckn ckn*} {**cak** [**encrypted** *encrypted_cak*] | *cak*} | **ports** [*port_list*] [**enable** | **disable**]] command using the **disable** option.

Example

The following example deletes the CA "testca":

```
# delete macsec connectivity-association testca
```

History

This command was first available in ExtremeXOS 30.1.

Platform Availability

This command is available on the following platforms.



Note

The MACsec feature requires the installation of the MAC Security feature pack license.

Platform	Ports
ExtremeSwitching 5320	All ports of all models except stacking ports.
ExtremeSwitching 5420	All ports of all models except stacking ports.
ExtremeSwitching 5520	All ports, except 5520-VIM-4X and 5520-24X 10G ports
ExtremeSwitching 5720	All ports of all models except stacking ports.

delete meter

```
delete meter meter-name
```

Description

Deletes a meter.

Syntax Description

<i>meter-name</i>	Specifies the meter name.
-------------------	---------------------------

Default

N/A.

Usage Guidelines

None.

Example

The following command deletes the meter `maximum_bandwidth`:

```
delete meter maximum_bandwidth
```

History

This command was available in ExtremeXOS 11.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

delete mirror name

```
delete mirror mirror_name {control_index} | all
```

Description

Deletes a user-defined mirroring instance, and unconfigures the "DefaultMirror" instance.

Syntax Description

<i>mirror_name</i>	Specifies a specific mirror name to delete.
<i>control_index</i>	Mirror destination control index (1-4). Also know as: etsysMirrorDestinationControllIndex. Each comprises a group of mirror names.
all	Specifies that you delete all named mirror instances.

Default

Disabled.

Usage Guidelines

Use this command to delete a user-defined mirroring instance and unconfigure the "DefaultMirror" instance. Mirroring instances must be in the "disabled" state in order to be deleted. The `all` command will fail if any mirroring instance is in the "enabled" state.

Example

The following example deletes all mirroring instances:

```
delete mirror all
```

History

This command was first available in ExtremeXOS 15.3.

Variable **control_index** to support policy-based mirrors was added in ExtremeXOS 30.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

delete mlag peer

```
delete mlag peer peer_name
```

Description

Deletes a peer switch from the structure.

Syntax Description

<i>peer_name</i>	Specifies an alpha numeric string identifying the MLAG peer.
------------------	--

Default

N/A.

Usage Guidelines

This command deletes an MLAG peer switch from the association structure.

Before you delete an MLAG peer switch, you must disable it. If it is not disabled, the following error message is displayed:

```
ERROR: MLAG ports currently associated with peer. First disable MLAG  
ports using "disable mlag port <port>" before deleting MLAG peer
```

Example

The following command deletes a peer switch structure switch101:

```
# delete mlag peer switch101
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

delete mpls rsvp-te lsp

```
delete mpls rsvp-te lsp [lsp_name | all]
```

Description

Deletes internal resources for the specified RSVP-TE LSP.

Syntax Description

<i>lsp_name</i>	Specifies the LSP within the switch to be deleted.
all	Deletes all RSVP-TE configured LSPs.

Default

N/A.

Usage Guidelines

This command deletes internal resources for the specified RSVP-TE LSP. The LSP is first withdrawn if it is currently active. Deleting an LSP may cause a PW to fail. Any static routes configured to a deleted LSP are also removed.

Example

The following command deletes the configured RSVP-TE LSP named lsp598:

```
delete mpls rsvp-te lsp lsp598
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

delete mpls rsvp-te path

```
delete mpls rsvp-te path [path_name | all]
```

Description

Deletes a configured RSVP-TE routed path with the specified path name.

Syntax Description

<i>path_name</i>	Specifies a path within the switch to be deleted.
all	Deletes all paths not associated with an LSP.

Default

N/A.

Usage Guidelines

This command deletes a configured RSVP-TE routed path with the specified name. All associated configuration information for the specified path is deleted. If the all keyword is specified, all paths not associated with an LSP are deleted.



Note

A path cannot be deleted as long as the path name is associated with an LSP.

Example

The following command deletes the configured RSVP-TE path named path598:

```
delete mpls rsvp-te path path598
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support *MPLS* as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

delete mpls rsvp-te profile

```
delete mpls rsvp-te profile [profile_name | all]
```

Description

Deletes a configured RSVP-TE profile with the specified profile name.

Syntax Description

<i>profile_name</i>	Specifies a configured RSVP-TE profile to be deleted.
all	Deletes all profiles not associated with an LSP, except the default profile.

Default

N/A.

Usage Guidelines

This command deletes a configured RSVP-TE profile with the specified profile name. If the all keyword is specified, all profiles not associated with an LSP are deleted (except for the default profile).



Note

A profile cannot be deleted as long as the profile name is associated with a configured LSP. The **default** profile cannot be deleted.

Example

The following command deletes the configured RSVP-TE profile named prof598:

```
delete mpls rsvp-te profile prof598
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

delete mpls static lsp

```
delete mpls static lsp [lsp_name | all]
```

Description

Deletes internal resources for one or all static LSPs.

Syntax Description

<i>lsp_name</i>	Identifies the LSP to be deleted.
all	Specifies that all LSPs are to be deleted.

Default

N/A.

Usage Guidelines

All resources associated with the specified LSPs are released. Static LSPs cannot be deleted when the LSP is configured for an IP route or VPLS configuration.

Example

The following command deletes a static LSP:

```
delete mpls static lsp lsp598
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

delete msdp mesh-group

```
delete msdp mesh-group mesh-group-name {vr vrname}
```

Description

Removes an [MSDP](#) mesh-group.

Syntax Description

<i>mesh-group-name</i>	Specifies the name of the MSDP mesh-group. The character string can be a maximum of 31 characters.
<i>vrname</i>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

N/A.

Usage Guidelines

A mesh-group is a group of MSDP peers with fully meshed MSDP connectivity. Mesh-groups are used to achieve two goals:

- Reduce SA message flooding.
- Simplify peer-RPF flooding.

SA messages received from a peer in a mesh-group are not forwarded to other peers in the same mesh-group.

Use the `delete msdp mesh-group` command only if you created a mesh-group that you want to remove. By default, there is no MSDP mesh-group.

Example

The following example removes a mesh-group called "verizon":

```
delete msdp mesh-group verizon
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

delete msdp peer

```
delete msdp peer [all | remoteaddr] {vr vr_name}
```

Description

Deletes an *MSDP* peer.

Syntax Description

all	Deletes all MSDP peers.
<i>remoteaddr</i>	Specifies the IP address of the MSDP router to configure as an MSDP peer.
<i>vr_name</i>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

N/A.

Usage Guidelines

None.

Example

The following example deletes an MSDP peer:

```
delete msdp peer 192.168.45.43
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

delete netlogin local-user

```
delete netlogin local-user user-name
```

Description

Deletes a specified local network login user name and its associated password.

Syntax Description

<i>user-name</i>	Specifies a local network login user name.
------------------	--

Default

N/A.

Usage Guidelines

Use the `show netlogin local-users` command to determine which local network login user name you want to delete from the system. The `show netlogin local-users` output displays the user name and password in a tabular format.

This command applies only to web-based and MAC-based modes of network login. 802.1X network login does not support local database authentication.

You must have administrator privileges to use this command.

Example

The following command deletes the local network login megtest along with its associated password:

```
delete netlogin local-user megtest
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

delete ntp key

```
delete ntp key [keyid | all]
```

Description

Deletes an NTP key; it cannot be used for outgoing or incoming NTP sessions.

Syntax Description

<i>keyid</i>	Specifies the key ID as a value from 1 to 65534.
all	Deletes all keys.

Default

N/A.

Usage Guidelines

N/A.

Example

The following command deletes NTP key 5 on the switch:

```
delete ntp key 5
```

The following command deletes all NTP keys on the switch:

```
delete ntp key all
```

History

This command was first available in ExtremeXOS 12.7.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

delete ospf area

```
delete ospf area [area-identifier | all]
```

Description

Deletes an *OSPF* area or all OSPF areas.

Syntax Description

<i>area-identifier</i>	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. Also, area 0.0.0.0 cannot be deleted.

Example

The following command deletes an OSPF area:

```
delete ospf area 1.2.3.4
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [*Switch Engine 32.3 Feature License Requirements*](#) document.

delete ospfv3 area

```
delete ospfv3 area [area_identifier | all]
```

Description

Deletes an OSPFv3 area or all OSPFv3 areas.

Syntax Description

<i>area_identifier</i>	Specifies an OSPFv3 area, a four-byte, dotted decimal number.
all	Specifies all areas.

Default

N/A.

Usage Guidelines

An OSPFv3 area cannot be deleted if it has an associated interface. Also, area 0.0.0.0 cannot be deleted.

Example

The following command deletes an OSPFv3 area:

```
delete ospfv3 area 1.2.3.4
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the Switch Engine 32.3 Feature License Requirements document.

delete policy access-list

```
delete policy access-list [all-rules | list_dot_rule]
```

Description

Deletes previously created access list and their rules.

Syntax Description

access-list	Configures access list rule model.
all-rules	Deletes all access lists and their rules.
<i>list_dot_rule</i>	Defines the access list name with optional rule name in the format <i>list_name.rule_name</i> .

Default

N/A.

Usage Guidelines

You can remove a specific rule or remove all the rules from an access list, or remove all access lists and their rules.

Example

The following example deletes the access list rule "ACL1.rule1":

```
# delete policy access-list ACL1.ace2
```

The following example deletes the access list "ACE":

```
# delete policy access-list ACE
```

The following example deletes all access lists and their rules:

```
# delete policy access-list all-rules
```

History

This command was first available in ExtremeXOS 30.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

delete policy access-list action-set

```
delete policy access-list action-set set-id
```

Description

Deletes a pre-defined action set.

Syntax Description

access-list	Specifies access-list features.
action-set	Specifies deleting an action set, which is a defined a set of actions that can be applied to multiple sets of match conditions.
<i>set-id</i>	Specifies which action set to delete by its global action set ID.

Default

N/A.

Example

The following example deletes action set "1":

```
# delete policy access-list action-set 1
```

History

This command was first available in ExtremeXOS 30.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

delete ports group

```
delete ports group port_group
```

Description

This command deletes a generic port-group name that can be associated with a list of ports. The *port_group* option could be implemented in configure or show commands that currently accept a *port_list*. The [QoS](#) commands are expanded to accept the *port_group* option. QoS commands that use port groups are updated automatically if the ports group is removed or if ports are added or removed from the group.

Syntax Description

<i>port_group</i>	Specifies a port group name.
-------------------	------------------------------

Default

N/A.

Usage Guidelines

Use this command to delete a generic port-group name associated with a list of ports.

Example

```
delete port-group testGroup
```

History

This command was first available in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

delete private-vlan

```
delete private-vlan name
```

Description

Deletes the PVLAN framework with the specified name.

Syntax Description

<i>name</i>	Specifies the name of the PVLAN to be deleted.
-------------	--

Default

N/A.

Usage Guidelines

The PVLAN is a framework that links network and subscriber VLANs; it is not an actual VLAN.

This command deletes the PVLAN framework, but it does not delete the associated VLANs. If the ports in the network VLAN were set to translate, they are changed to tagged.

Example

The following example deletes the PVLAN named "companyx":

```
delete private-vlan companyx
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms that support the Private VLAN feature. For features and the platforms that support them, see the [Switch Engine 32.3 Feature License Requirements](#) document.

delete process

delete process

Description

This command provides the ability for an end-user to delete dynamically-created processes.

Syntax Description

This command has no arguments or variables.

Default

N/A

Usage Guidelines

Use this command to delete dynamically-created processes only.

History

This command was first available in ExtremeXOS 15.7..

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

delete protocol

delete protocol {filter} *filter_name*

Description

Deletes a user-defined protocol.

Syntax Description

filter	Deletes a protocol filter.
<i>filter_name</i>	Specifies a protocol filter name to delete.

Default

N/A.

Usage Guidelines

If you delete a protocol that is in use by a VLAN, the protocol associated with than VLAN becomes none.

Example

The following examples delete a protocol named "my_filter" and a protocol filter named "my_other_filter":

```
delete protocol "my_filter"
delete protocol filter "my_other_filter"
```

History

This command was first available in ExtremeXOS 10.1.
The **filter** keyword was added in ExtremeXOS 15.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

delete qosprofile

```
delete qosprofile [QP2 | QP3 | QP4 | QP5 | QP6 | QP7]
```

Description

Deletes a user-created QoS profile.

Syntax Description

QP2 . . . QP7	Specifies the user-created QoS profile you want to delete.
---------------	--

Default

N/A.

Usage Guidelines

You cannot delete the default QoS profiles of QP1 and QP8. On a SummitStack, you also cannot delete QoS profile QP7. If you attempt to delete QoS profile QP7, the system returns an error.

All configuration information associated with the specified QoS profile is removed.

Example

The following command deletes the user-created QoS profile QP3:

```
delete qosprofile qp3
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

delete sshd2 user-key

```
delete sshd2 user-key key_name
```

Description

Deletes a user key.

Syntax Description

<i>key_name</i>	Specifies the name of the public key to be deleted.
-----------------	---

Default

N/A.

Usage Guidelines

This command is used to delete a user key. The key is deleted regardless of whether or not it is bound to a user.



Note

If a user is bound to the key, they are first unbound or unassociated, and then the key is deleted.

Example

The following example shows the SSH user key id_dsa_2048 being deleted:

```
delete sshd2 user-key id_dsa_2048
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

delete stpd

```
delete stpd stpd_name
```

Description

Removes a user-defined *STPD* from the switch.

Syntax Description

<i>stpd_name</i>	Specifies a user-defined STPD name on the switch.
------------------	---

Default

N/A.

Usage Guidelines

If your STPD has the same name as another component, for example a *VLAN*, we recommend that you specify the identifying keyword as well as the name. If you do not specify the stpd keyword, an error message similar to the following is displayed:

```
%% Ambiguous command: "delete Test"
```

In this example, to delete the STPD Test, enter delete stpd Test.

If you created an STPD with a name unique only to that STPD, the keyword stpd is optional.

The default STPD, s0, cannot be deleted.

In an *MSTP* environment, you cannot delete or disable a CIST if any of the MSTIs are active in the system.

Example

The following example deletes an STPD named "purple_st":

```
delete stpd purple_st
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

delete tunnel

```
delete tunnel tunnel_name
```

Description

Deletes an IPv6 tunnel.

Syntax Description

<i>tunnel_name</i>	Specifies an IPv6 tunnel.
--------------------	---------------------------

Default

N/A.

Usage Guidelines

This command will destroy a previously created tunnel. The command acts on either a 6to4 or a 6in4 tunnel. When the tunnel interface is removed, all dynamic routes through that interface are purged from the system. The configured static routes are removed from the hardware tables and become inactive.

Example

The following example deletes the tunnel link39:

```
delete tunnel link39
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 interworking feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

delete upm profile

```
delete upm profile profile-name
```

Description

Deletes the specified profile.

Syntax Description

<i>profile-name</i>	Specifies the UPM profile to be deleted.
---------------------	--

Default

N/A.

Example

The following command deletes a UPM profile called sample_1:

```
delete upm profile sample_1
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the Universal Port feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

delete upm timer

```
delete upm timer timer-name
```

Description

Deletes the specified UPM timer.

Syntax Description

<i>timer-name</i>	Specifies the name of the UPM timer to be deleted.
-------------------	--

Default

N/A.

Usage Guidelines

You can delete a UPM timer by specifying its name.

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the Universal Port feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

delete var

`delete var varname`



Note

This is a script command and operates only in scripts or on the command line when scripting is enabled with the following command: `enable cli scripting {permanent}`.

Description

Deletes a variable.

Syntax Description

<code>varname</code>	Specifies the name of the scripting variable to be deleted.
----------------------	---

Default

N/A.

Usage Guidelines

The format of a local variable (case insensitive) is: \$VARNAME.

Example

The following example deletes local variable x:

```
delete var x
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

delete var key

delete var key *key*



Note

This is a script command and operates only in scripts or on the command line when scripting is enabled with the following command: `enable cli scripting {permanent}`.

Description

Deletes the variables that have been saved using a key.

Syntax Description

<i>key</i>	Specifies that variables associated with the specified key must be deleted.
------------	---

Default

N/A.

Usage Guidelines

CLI scripting must be enabled to use this command. The user is responsible for generating unique keys for each variable. The system has a limited amount of memory to store these variables.

Example

The following command deletes all variables associated with the key “red:”

```
delete var key red
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

delete virtual-network

delete **virtual-network** *vn_name*

Description

This command deletes a virtual network.

Syntax Description

virtual-network	Designates deleting a virtual network.
<i>vn_name</i>	Specifies which virtual network.

Default

N/A.

Example

The following example deletes the virtual network "my_virtual_network":

```
delete virtual-network my_virtual_network
```

History

This command was first available in ExtremeXOS 21.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

delete virtual-network remote-endpoint vxlan ipaddress

```
delete virtual-network remote-endpoint vxlan ipaddress ipaddress {vr  
  vr_name}
```

Description

This command deletes a remote endpoint.

Syntax Description

<i>ipaddress</i>	A remote endpoint IP address.
vr	VR/VRF instance the remote endpoint is associated with
<i>vr_name</i>	An existing VR/VRF name.

Default

N/A.

Usage Guidelines

This command is useful when user wants to delete a remote endpoint in addition to the ones learned dynamically (*OSPF* extensions).

Example

To remove a remote endpoint:

```
delete virtual-network vxlan remote-endpoint ipaddress 1.2.3.4
```

History

This command was first available in ExtremeXOS 21.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

delete virtual-router

```
delete virtual-router vr-name
```

Description

This command deletes a VR or VRF.

Syntax Description

<i>vr-name</i>	Specifies the name of the VR or VRF.
----------------	--------------------------------------

Default

N/A.

Usage Guidelines

Only user VRs and VRFs can be deleted.

Before you delete a user VR, you must delete all VLANs and protocols assigned to the VR, and you must delete any child VRFs. All of the ports assigned to a deleted VR are made available to assign to other VRs.

Before you delete a VRF, you must delete all VLANs and stop all protocols that are assigned to that VRF. All of the ports assigned to a deleted VRF are deleted and made available to assign to other VRs and VRFs. Any routing protocol instance that is assigned to the VRF is deleted gracefully.

Example

The following example deletes the VR "vr-acme":

```
delete virtual-router vr-acme
```

The following example deletes the VRF "vrf1":

```
delete virtual-router vrf1
```

History

This command was first available in ExtremeXOS 11.0.

Support for non-VPNVRFs was added in ExtremeXOS 12.5.

Support for VPN VRFs was added in ExtremeXOS 12.6.0-BGP.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

delete vlan

```
delete [ {vlan} vlan_name | vlan vlan_list]
```

Description

Deletes a *VLAN*.

Syntax Description

<i>vlan_name</i>	Specifies a VLAN name.
<i>vlan_list</i>	Specifies a VLAN list of IDs.

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 `configure svlan delete ports` command.



Note

The default VLAN cannot be deleted. Before deleting an ISC VLAN, you must delete the peer.

Example

The following command deletes the VLAN accounting:

```
delete accounting
```

History

This command was first available in ExtremeXOS 10.1.

The *vlan_list* option was added in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

delete vman

```
delete vman [vman-name | vman_list]
```

Description

Deletes a previously created VMAN.

Syntax Description

<i>vman-name</i>	Specifies a VMAN name.
<i>vman_list</i>	Specifies a VMAN list.

Default

N/A.

Usage Guidelines

None.

Example

The following command deletes the VMAN accounting:

```
delete vman accounting
```

History

This command was first available in ExtremeXOS 11.0.

The *vman_list* variable was added in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

delete vm

```
delete vm vm_name
```

Description

Deletes an existing virtual machine (VM).

Syntax Description

vm	Designates a virtual machine.
<i>vm_name</i>	Specifies the VM name to delete.

Default

N/A.

Usage Guidelines

The Integrated Application Hosting (IAH) feature requires the Solid State Storage Device SSD-120.

Example

The following example deletes the vm "vm1":

```
# delete vm vm1
```

History

This command was first available in ExtremeXOS 30.3.

Stop the VM before attempting to delete it (`stop vm vm_name [forceful | graceful]`).

Platform Availability

This command is available on all platforms that support the Integrated Application Hosting (IAH) feature and have a Core or Premier license installed. For a list of platforms that support the IAH feature and for information about licenses, see the [Switch Engine 32.3 Feature License Requirements](#).

delete vm-tracking local-vm

```
delete vm-tracking local-vm {mac-address mac}
```

Description

Deletes the specified VM entry in the local VM database.

Syntax Description

<i>mac</i>	Specifies the MAC address for a VM entry to delete.
------------	---

Default

N/A.

Usage Guidelines

None.

Example

The following command deletes the VM entry for MAC address 00:E0:2B:12:34:56 in the local VM database:

```
# delete vm-tracking local-vm mac-address 00:E0:2B:12:34:56
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

delete vm-tracking vpp

```
delete vm-tracking vpp {vpp_name}
```

Description

Deletes the specified LVPP.

Syntax Description

<i>vpp_name</i>	Specifies a name for the LVPP to delete.
-----------------	--

Default

N/A.

Usage Guidelines

None.

Example

The following command deletes the VPP named vpp1:

```
# delete vm-tracking vpp vpp1
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

delete vpls

```
delete vpls [vpls_name | all]
```



Note

This command has been replaced with the following command: `delete l2vpn [vpls [vpls_name | all] | vpws [vpws_name | all]]`.

This command is still supported for backward compatibility, but it will be removed from a future release, so we recommend that you start using the new command.

Description

Deletes the VPLS with the specified *vpls_name*.

Syntax Description

<i>vpls_name</i>	Identifies the VPLS within the switch (character string).
all	Specifies all VPLS.

Default

N/A.

Usage Guidelines

This command deletes the VPLS with the specified *vpls_name*. All PWs established to VPLS peers are terminated. The **all** keyword may be used to indicate that all VPLS instances are to be deleted.

Example

This commands deletes the VPLS myvpls:

```
delete vpls myvpls
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support *MPLS* as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

delete vrrp group

```
delete vrrp group group_name
```

Description

This command deletes a *VRRP* group used to operate in high-scale mode.

Syntax Description

group	Specifies deleting a VRRP group.
<i>group_name</i>	Specifies the VRRP group name.

Default

None.

Example

The following example deletes a VRRP group called "vrrp1".

```
delete vrrp group vrrp1
```

History

This command was first available in ExtremeXOS 22.2.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the VRRP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

delete vrrp vlan vrid

```
delete vrrp vlan vlan_name vrid vridval
```

Description

Deletes a specified [VRRP](#) instance.

Syntax Description

<i>vlan_name</i>	Specifies the name of a VRRP VLAN .
<i>vridval</i>	Specifies the VRID for the VRRP instance. To display the configured VRRP router instances, enter the show vrrp command.

Default

N/A.

Usage Guidelines

None.

Example

The following command deletes the VRRP instance on the VLAN vrrp-1 identified by VRID 2:

```
delete vrrp vlan vrrp-1 vrid 2
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the VRRP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

delete xml-notification target

```
delete xml-notification target target
```

Description

Deletes the Web server target on the XML client process.

Syntax Description

<i>target</i>	Specifies the configured target.
---------------	----------------------------------

Default

N/A.

Usage Guidelines

Use this command to delete the Web server target on the XML client process.

Example

The following command deletes the target test2:

```
delete xml-notification target test2
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable access-list permit to-cpu

```
disable access-list permit to-cpu
```

Description

Allows special packets to be blocked by low priority [ACLs](#).

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

This command allows ACLs to deny certain special packets from reaching the CPU, even if the packets match ACLs that would otherwise deny them. The special packets include [STP](#) and [EAPS](#) BPDUs, and ARP replies for the switch.

When this feature is disabled, these same packets will be denied if an ACL is applied that contains a matching entry that denies the packets. Contrary to expectations, the packets will still be denied if there is a higher precedence entry that permits the packets.

To enable this feature, use the following command:

```
enable access-list permit to-cpu
```

Example

The following example enables ACLs to deny STP BPDU packets from reaching the switch CPU:

```
disable access-list permit to-cpu
```

History

This command was first available in ExtremeXOS 11.3.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable access-list refresh blackhole

```
disable access-list refresh blackhole
```

Description

Disables blackholing of packets during [ACL](#) refresh.

Syntax Description

This command has no arguments or variables.

Default

The feature is enabled.

Usage Guidelines

When access control lists (ACLs) are refreshed, this feature provides that any packets arriving during the refresh will be blackholed.

If you disable this feature, the ACLs will be refreshed as described in the [refresh policy](#) command.

To enable this feature, use the following command:

```
enable access-list refresh blackhole
```

Example

The following command disables dropping of packets during an ACL refresh:

```
disable access-list refresh blackhole
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable account

```
disable account [all {admin|user | name]
```

Description

Disables the specified account locally.

Syntax Description

all	Specifies that all accounts, or all accounts of a certain type, will be disable.
admin	Specifies that all administrative accounts will be disabled locally.
user	Specifies that all user accounts, including Lawful-Intercept accounts, will be disabled locally.
<i>name</i>	Specifies the name of the account that will be disabled locally.

Default

Enabled.

Usage Guidelines

If the user is disabled locally, the user's login will fail.

Disabling accounts affects the following northbound interfaces:

- Console
- TELNET
- SSH
- HTTP
- XML

If you disable all administrative accounts, you can use the failsafe account.

Example

The following example disables all user accounts.

```
disable account all user
```

History

This command was first available in ExtremeXOS 15.7.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable auto-provision

```
disable auto-provision
```

Description

Disables the auto provision capability.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to disable the auto provision capability.

To display the status of auto provision on the switch, use the [show auto-provision](#) command.

Example

The following command disables the auto provision capability:

```
disable auto-provision
```

The following message is displayed:

```
# disable auto-provision
This setting will take effect at the next reboot of this switch.
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable avb

```
disable avb
```

Description

This command is a macro command that can be used to disable all AVB protocols globally on the switch. It is equivalent to issuing the following three commands:

```
disable mvrp
```

```
disable msrp
```

```
disable network-clock gptp
```

Syntax Description

avb	Audio Video Bridging
------------	----------------------

Default

Disabled.

Usage Guidelines

Use this command to disable all AVB protocols globally on the switch.

Example

```
disable avb
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable avb ports

```
disable avb ports [port_list | all]
```

Description

This command is a macro command that can be used to disable all AVB protocols on the given ports. It is equivalent to issuing the following three commands:

```
disable mvrp ports [port_list | all]
```

```
disable msrp ports [port_list | all]
```

```
disable network-clock gptp ports [port_list | all]
```

Syntax Description

avb	Audio Video Bridging.
<i>port_list</i>	Port list separated by a comma or "-".
all	All ports.

Default

Disabled.

Usage Guidelines

Use this command to disable all AVB protocols on the given ports.

Example

```
disable avb ports all
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable bgp

```
disable bgp
```

Description

Disables *BGP*.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

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 ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

disable bgp advertise-inactive-route

```
disable bgp {address-family [ipv4-unicast | ipv4-multicast | ipv6-unicast | ipv6-multicast]} advertise-inactive-route
```

Description

Disables advertisement of *BGP* inactive routes, which are defined as those routes that rated best by BGP and not best in the IP routing table.

Syntax Description

address-family	Specifies an IPv4 or IPv6 unicast or multicast address family.
-----------------------	--

Default

Disabled.

If no address family is specified, IPv4 unicast is the default address family.

Usage Guidelines

This command can be successfully executed only when BGP is globally disabled. If you want to disable inactive route advertisement and BGP is enabled, you must disable BGP ([disable bgp](#)), disable this feature, and then enable BGP ([enable bgp](#)).

If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.

Example

The following command disables inactive route advertisement for IPv4 unicast traffic:

```
disable bgp address-family ipv4-unicast advertise-inactive-route
```

History

This command was first available in ExtremeXOS 12.1.3.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

disable bgp aggregation

```
disable bgp aggregation
```

Description

Disables BGP route aggregation.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

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 ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

disable bgp always-compare-med

```
disable bgp always-compare-med
```

Description

Disables BGP from comparing Multi Exit Discriminators (MEDs) for paths from neighbors in different Autonomous Systems (AS).

Syntax Description

This command has no arguments or variables.

Default

ExtremeXOS does not compare MEDs for paths from neighbors in different AS.

Usage Guidelines

The MED is one of the parameters that is considered when selecting the best path among many alternative paths. The path with a lower MED is preferred over a path with a higher MED. By default, during the best path selection process, MED comparison is done only among paths from the same AS.

BGP must be disabled before you can change the configuration with this command.

Example

The following command disables MED from being used in comparison among paths from different AS:

```
disable bgp always-compare-med
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

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 commands.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

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 AS-number : number
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the *BGP* feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

disable bgp export vr

```
disable bgp export {vr} vr_name route_type {address-family} vpnv4
```

Description

For IPv4 and IPv6 routes, this command disables the PE router to export and redistribute local VRF routes to remote PE routers through *BGP*.

Syntax Description

vr	Specifies the source VPN VRF of the exported routes.
<i>vr_name</i>	Specifies the name of the source VPN VRF.
<i>route_type</i>	Specifies the source or origin of the route types to be exported to remote PE routers. Valid Types: blackhole, direct, and bgp.
<i>address-family</i>	Specifies the address family for the exported routes. Valid types are ipv4-unicast, vpnv4.
vpnv4	Specifies that routes from the VRF are exported as vpnv4 routes over MPBGP.

Default

Disabled.

Usage Guidelines

This command disables a PE router to advertise learned routes from CE routers to remote PE routers in a Service Provider's backbone. Executing this command allows the PE router to convert VRF native IPv4 routes into VPN-IPv4 routes and advertise to all remote PE BGP neighbors as VPN-IPv4 routes.

- For Layer 3 VPNs, you must enter the `disable bgp export vr` command in the context of the VRF that supports the Layer 3 VPN.
- When the export source is the Layer 3 VPN, you can specify `direct`, or `remote-vpn` to disable route export to the VRF. The destination address family must be `ipv4-unicast`.
- This export command is applicable in Parent VR context only. If you execute it in a VRF context, an error message is returned.

- The source VPN VRF must be a child of the Parent VR.
- BGP need not be added to a VPN VRF to export routes from a VPN VRF.
- The direction of where the redistribution is targeted is implicit on the keywords used, For eg:- remote-vpn only applies to remote routes from PE redistributed to CE, hence we cannot use it with address family vpnv4. Similarly bgp only applies to EBGp routes from CE exported as VPN routes, hence we use it only with address family vpnv4. Other sources such as “static” and “direct” are redistributed both ways.

Example

The following command disables BGP to advertise a vpnv4 route named "corp1_vpn_vrf":

```
disable bgp export "corp1_vpn_vrf" bgp address-family vpnv4
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

The blackhole option was added in ExtremeXOS 12.1.3.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Support for Layer 3 VPNs was added in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

disable bgp export

```
disable bgp export route_type {{address-family} address_family}
```

For Layer 3 VPNs:

```
disable bgp export route_type {{address-family} address_family}
```

Description

Disables *BGP* from exporting routes from other protocols to BGP peers.

Syntax Description

bgp	For Layer 3 VPNs, this specifies that BGP routes learned from CE routers are to be exported to remote PE routers.
<i>route_type</i>	Specifies the BGP export route type.
address-family	Specifies an IPv4 or IPv6 unicast or multicast address family.

Default

Disabled.

If no address family is specified, IPv4 unicast is the default.



Note

You must specify an IPv6 address family for an IPv6 peer, because an IPv6 peer does not support the default IPv4 unicast address family. Similarly, if you specify an IPv4 peer and an address family in the command, an IPv4 address family must be specified.

Usage Guidelines

The exporting of routes between any two routing protocols is a discrete 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 policies to associate BGP attributes including Community, NextHop, MED, Origin, and Local Preference with the routes. Policies can also be used to filter out exported routes.

Using the export command to redistribute routes complements the redistribution of routes using the `configure bgp add network` command. The `configure 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.



Note

For this command to execute, the specified protocol must support the specified address family. For example, the command fails if you specify OSPF and the IPv6 unicast address family. You can specify blackhole, direct, static, and IS-IS routes with IPv4 or IPv6 address families.

For Layer 3 VPNs, the disable bgp export command must be entered in the context of the VRF that supports the Layer 3 VPN.

When the export source is the Layer 3 VPN, you can specify direct, or remote-vpn to disable route export to the VRF. The destination address family must be ipv4-unicast.

When the export source is the VRF, you can specify `direct`, or `bgp` to disable route export to the VPN. The destination address family must be `vpn4`.

Example

The following command disables BGP from exporting routes from the OSPF protocol to BGP peers:

```
disable bgp export ospf
```

The following command disables the export of BGP routes from a VRF to a VPN:

```
disable bgp export bgp address-family vpnv4
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

The blackhole option was added in ExtremeXOS 12.1.3.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Support for Layer 3 VPNs was added in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

disable bgp export [static | direct] l2vpn-evpn

```
disable bgp export [static | direct] {address-family address_family}  
      l2vpn-evpn {vr vr_name}
```

Description

Disables export of direct, static, and BGP routes from a VRF into BGP, running on the specified VR, as EVPN routes to be advertised by BGP as Type 5 routes.

Syntax Description

bgp	Specifies showing the BGP configuration.
export	Specifies redistributing information from another routing protocol.
static	Specifies static routes.
direct	Specifies direct routes.

address-family	Specifies the address family.
<i>address_family</i>	Sets the address family type.
l2vpn-evpn	Specifies the L2VPN EVPN address family.
vr	Specifies the source VR.
<i>vr_name</i>	Designates the source VR name. Both VPN-VRFs and non-VPN-VRFs are supported.

Default

N/A.

Usage Guidelines

N/A.

Example

The following example disables exports of static routes on VR "vr-a" as EVPN routes to be advertised by BGP as Type 5 routes:

```
# disable bgp export static l2vpn-evpn vr vr-a
```

History

This command was first available in ExtremeXOS 30.6.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

disable bgp fast-external-fallover

```
disable bgp fast-external-fallover
```

Description

Disables BGP fast external fallover functionality.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command disables the BGP fast external fallover on the router. This command applies to all directly-connected external BGP neighbors.

When BGP fast external fallover is enabled, the directly-connected EBGP neighbor session is immediately reset when the connecting link goes down.

If BGP fast external fallover is disabled, BGP waits until the default hold timer expires (3 keepalives) to reset the neighboring session. In addition, BGP might teardown the session somewhat earlier than hold timer expiry if BGP detects that the TCP session and it's directly connected link is broken (BGP detects this while sending or receiving data from TCP socket).

Example

The following command disables BGP fast external fallover:

```
disable bgp fast-external-fallover
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

disable bgp mpls-next-hop

```
disable bgp mpls-next-hop
```

Description

Disables IP forwarding over calculated MPLS LSPs to subnets learned via BGP.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command disables IP forwarding over calculated MPLS LSPs to subnets learned via BGP. (Calculated refers to an LSP that only reaches part of the way to the destination). By default, IP forwarding over MPLS LSPs to subnets learned via BGP is disabled.

Example

The following command disables BGP's use of MPLS LSPs to reach BGP routes:

```
disable bgp mpls-next-hop
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

disable bgp multipath-relax

```
disable bgp multipath-relax
```

Description

Disables BGP multipath-relax feature, which modifies the definition of an equal cost BGP route.

Syntax Description

multipath-relax	Selects BGP multipath relax feature.
------------------------	--------------------------------------

Default

This feature is disabled by default.

Usage Guidelines

This feature modifies the definition of equal cost BGP routes as specified in *RFC-4271*. In particular, routes with the same AS-path length, but differing AS numbers in the path are not considered equal cost by default. However, with multipath-relax enabled, routes with the same AS-path length can have differing AS number values in the AS-path and still be considered equal cost.

BGP must be disabled (`disable bgp`) first to disable this feature.

Example

The following example disables the BGP multipath-relax feature:

```
disable bgp multipath-relax
```

History

This command was first available in ExtremeXOS 22.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

disable bgp neighbor address-family l2vpn-evpn

```
disable bgp {neighbor [remoteaddr | all]} {{address-family} l2vpn-evpn}  
next-hop-unchanged
```

Description

Disables overriding the *BGP* specification behavior with respect to the next-hop of routes advertised to EBGp peers.

Syntax Description

bgp	Specifies BGP.
neighbor	Specifies BGP neighbor.
<i>remoteaddr</i>	Specifies BGP neighbor IP address.
all	Specifies all BGP neighbors.
address-family	Specifies address family.
l2vpn-evpn	Specifies L2VPN EVPN address-family type.
next-hop-unchanged	Enables preserving the BGP next-hop when routes are advertised to EBGp peers (default is disabled).

Default

Default is that next-hop-unchanged is disabled.

Usage Guidelines

This command disables overriding the specification behavior with respect to the next-hop of routes advertised to EBGp peers. Specifically, disabling with this command does not maintain the BGP next-hop for routes advertised to EBGp peers instead of replacing the next-hop with either the outgoing interface IP address or the local loopback address.

Example

The following example disables next-hop unchanged for BGP neighbor at 192.168.66.2:

```
# disable bgp neighbor 192.168.66.2 l2vpn-evpn next-hop-unchanged
```

History

This command was first available in ExtremeXOS 30.2.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

disable bgp neighbor capability address-family vpnv4

```
disable bgp {neighbor} [all | remoteaddr] capability address-family  
    vpnv4 type [community | ext-community | prefix] {[send | receive |  
    both] }
```

Description

This command disables neighbor capability for one or all *BGP* neighbors on a Layer 3 VPN.

Syntax Description

all	Specifies that the configuration applies to all neighbors in the specified address family. If no address family is specified or if an IPv4 address is specified, the configuration applies to all IPv4 neighbors.
<i>remoteaddr</i>	Specifies the IPv4 address of a BGP neighbor.
vpnv4	Specifies the VPNv4 address family for Layer 3 VPN support.
community	Disables neighbor capability for communities.
ext-community	Disables neighbor capability for extended communities.
prefix	Disables neighbor capability for prefixes.
send	Disables neighbor capability filter list send capability.
receive	Disables neighbor capability filter list receive capability.
both	Disables neighbor capability filter list send and receive capability.

Default

Disabled.

If the direction is not specified, the both option applies.

Usage Guidelines

Enter this command multiple times to configure the address family, type, and direction attributes.

Example

The following command disables the neighbor capability feature for a Layer 3 VPN neighbor:

```
disable bgp neighbor 1.1.1.1 capability address-family vpnv4
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

disable bgp neighbor capability

```
disable bgp neighbor [all | remoteaddr] capability [ipv4-unicast | ipv4-multicast | ipv6-unicast | ipv6-multicast | vpnv4 | route-refresh | ipv4-vxlan | l2vpn-evpn]
```

Description

This command disables an address family or the route refresh capability for one or all neighbors.

Syntax Description

all	Specifies that the configuration applies to all neighbors in the specified address family. If no address family is specified or if an IPv4 address is specified, the configuration applies to all IPv4 neighbors. If an IPv6 address family is specified, the configuration applies to all IPv6 neighbors.
<i>remoteaddr</i>	Specifies the IPv4 or IPv6 address of a <i>BGP</i> neighbor. The switch uses the IP address format to determine if the address is an IPv4 or IPv6 address.
ipv4-unicast	Specifies the IPv4 unicast address family.
ipv4-multicast	Specifies an IPv4 multicast address family.
ipv6-unicast	Specifies the IPv6 unicast address family.
ipv6-multicast	Specifies an IPv6 multicast address family.
vpnv4	Specifies the VPNv4 address family for Layer 3 VPN support.
route-refresh	Specifies ROUTE-REFRESH message capabilities.

ipv4-vxlan	Specifies IPv4 VXLAN capability.
l2vpn-evpn	Specifies L2 VPN EVPN address family.

Default

The following capabilities are enabled by default for IPv4 peers: IPv4 unicast, IPv4 multicast, and route refresh.

The following capabilities are enabled by default for IPv6 peers: route refresh.

Usage Guidelines

This command applies to the current VR or VRF context.



Note

To inter-operate with Cisco routers for BGP graceful restart, you must enable IPv4 unicast address capability.

For an IPv6 peer, an IPv6 address family must be specified, because an IPv6 peer does not support IPv4 address families. If no address family is specified for an IPv6 peer, the default, the IPv4 unicast address family, applies and the command fails. Similarly an IPv4 peer only supports IPv4 address families and the command fails if an IPv6 address family is specified.

Example

The following example disables the route-refresh feature for all neighbors:

```
disable bgp neighbor all capability route-refresh
```

The following example disables the VPNv4 address family for a neighbor:

```
disable bgp neighbor 192.168.96.235 capability vpnv4
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Support for Layer 3 VPNs was added in ExtremeXOS 15.3.

Support for IPv4 VXLAN was added in ExtremeXOS 22.3.

Support for L2 VPN EVPN address family was added in ExtremeXOS.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

disable bgp neighbor originate-default

```
disable bgp [{neighbor} remoteaddr | neighbor all] {address-family  
[ipv4-unicast | ipv4-multicast | ipv6-unicast | ipv6-multicast] }  
originate-default
```

Description

Removes a default route to a single *BGP* neighbor or to all BGP neighbors.

Syntax Description

<i>remoteaddr</i>	Specifies the IPv4 or IPv6 address of a BGP neighbor.
all	Specifies that the configuration applies to all neighbors in the specified address family. If no address family is specified, the configuration applies to the IP Unicast family on all IPv4 peers. If an IPv4 address family is specified, the configuration applies to all IPv4 neighbors. If an IPv6 address family is specified, the configuration applies to all IPv6 neighbors.
address-family	Specifies an IPv4 or IPv6 unicast or multicast address family.

Default

Disabled. BGP does not automatically originate and advertise default routes to BGP neighbors.

If no address family is specified and an IPv4 address is detected, IPv4 unicast is the default address family.



Note

You must specify an IPv6 address family for an IPv6 peer, because an IPv6 peer does not support the default IPv4 unicast address family. Similarly, if you specify an IPv4 peer and an address family in the command, an IPv4 address family must be specified.

Usage Guidelines

This command can be successfully executed at any time, irrespective of whether local BGP or the remote BGP peer is enabled or disabled.

If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.

Example

The following command removes default routes for IPv4 unicast traffic for all BGP peer nodes:

```
disable bgp neighbor all originate-default
```

History

This command was first available in ExtremeXOS 12.1.3.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

disable bgp neighbor remove-private-AS-numbers

```
disable bgp neighbor [remoteaddr | all] remove-private-AS-numbers
```

Description

Disables the removal of private AS numbers from the AS path in route updates sent to EBGPs.

Syntax Description

<i>remoteaddr</i>	Specifies the IPv4 or IPv6 address of a BGP neighbor.
all	Specifies all IPv4 and IPv6 neighbors.

Default

Disabled.

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 private AS number can be stripped out from the AS paths of the advertised routes using this feature.

This command applies to the current VR or VRF context.

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 192.168.1.17 remove-private-as-numbers
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

disable bgp neighbor soft-in-reset

```
disable bgp neighbor [all | remoteaddr] {address-family [ipv4-unicast |  
ipv4-multicast | ipv6-unicast | ipv6-multicast | vpn4] } soft-in-reset
```

Description

Disables the soft input reset feature.

Syntax Description

all	Specifies that the configuration applies to all neighbors in the specified address family. If no address family is specified, the configuration applies to the IP Unicast family on all IPv4 peers. If an IPv4 address family is specified, the configuration applies to all IPv4 neighbors. If an IPv6 address family is specified, the configuration applies to all IPv6 neighbors.
<i>remoteaddr</i>	Specifies the IPv4 or IPv6 address of a BGP neighbor.
ipv4-unicast	Specifies the IPv4 unicast address family.
ipv4-multicast	Specifies an IPv4 multicast address family.
ipv6-unicast	Specifies the IPv6 unicast address family.
ipv6-multicast	Specifies an IPv6 multicast address family.
vpn4	Specifies the VPNv4 address family for Layer 3 VPN support.

Default

Disabled.

If no address family is specified and an IPv4 address is detected, IPv4 unicast is the default address family.

Usage Guidelines

Disabling the soft input reset feature can potentially limit the amount of system memory consumed by the RIB-in.

If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.



Note

For an IPv6 peer, an IPv6 address family must be specified, because an IPv6 peer does not support IPv4 address families. If no address family is specified for an IPv6 peer, the default, the IPv4 unicast address family, applies and the command fails. Similarly an IPv4 peer only supports IPv4 address families and the command fails if an IPv6 address family is specified.

Before you can change the configuration with this command, you must disable BGP, and you must disable the corresponding BGP neighbor session using the following command:

```
disable bgp neighbor [remoteaddr | all]
```

To disable this feature on Layer 3 VPNs, you must do so in the context of the *MPLS*-enabled VR; this feature is not supported for BGP neighbors on the CE (VRF) side of the PE router.

This command applies to the current VR or VRF context.

Example

The following command disables the soft input reset for the neighbor at 192.168.1.17:

```
disable bgp neighbor 192.168.1.17 soft-in-reset
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Support for Layer 3 VPNs was added in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

disable bgp neighbor

```
disable bgp neighbor [remoteaddr | all]
```

Description

Disables the *BGP* session.

Syntax Description

<i>remoteaddr</i>	Specifies the IPv4 or IPv6 address of a BGP neighbor.
all	Specifies all IPv4 and IPv6 neighbors.

Default

Disabled.

Usage Guidelines

After the session has been disabled, all the information in the route information base (RIB) for the neighbor is flushed.

This command applies to the current VR or VRF context.

Example

The following command disables the BGP session:

```
disable bgp neighbor 192.168.1.17
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

disable bgp peer-group capability address-family vpnv4

```
disable bgp peer-group peer-group-name capability address-family vpnv4
type [community | ext-community] {[send | receive | both]}
```

Description

This command disables peer-group capability for a peer group on a Layer 3 VPN.

Syntax Description

<i>remoteaddr</i>	Specifies the IPv4 address of a <i>BGP</i> neighbor.
vpnv4	Specifies the VPNv4 address family for Layer 3 VPN support.
community	Disables peer-group capability for communities.
ext-community	Disables peer-group capability for extended communities.
send	Disables peer-group capability filter list send capability.
receive	Disables peer-group capability filter list receive capability.
both	Disables peer-group capability filter list send and receive capability.

Default

Disabled.

If the direction is not specified, the both option applies.

Usage Guidelines

Enter this command multiple times to configure the address family, type, and direction attributes.

By specifying the address-family, type and direction in multiple commands, you can better control the actual ORF capabilities sent to a peer. In the case where a particular address-family is explicitly disabled for a peering, the ORF capability configuration for that address-family is ignored and not sent.

ORF capabilities can only be enabled for IPv4 neighbors, and only for IPv4 address families. If configured for IPv6 neighbors or address-families the command is rejected with the following error message:

```
Outbound-route-filtering not supported for IPv6 neighbors
```

or

```
Outbound-route-filtering not supported for address family <addr_family>
```

Example

The following command disables the peer-group capability feature for a Layer 3 VPN peer group:

```
disable bgp peer-group vpn capability address-family vpnv4
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

disable bgp peer-group capability

```
disable bgp peer-group peer-group-name capability [ipv4-unicast | ipv4-multicast | ipv6-unicast | ipv6-multicast | vpnv4 | route-refresh]
```

Description

This command disables an address family or the route-refresh capability for a peer group.

Syntax Description

<i>peer-group-name</i>	Specifies a peer group.
ipv4-unicast	Specifies the IPv4 unicast address family.
ipv4-multicast	Specifies an IPv4 multicast address family.
ipv6-unicast	Specifies the IPv6 unicast address family.
ipv6-multicast	Specifies an IPv6 multicast address family.
vpnv4	Specifies the VPNv4 address family for Layer 3 VPN support.
route-refresh	Specifies ROUTE-REFRESH message capabilities.

Default

All capabilities are enabled for IPv4 peer groups by default.

Only the route refresh capability is enabled for peer groups by default.

Usage Guidelines

This command applies to the current VR or VRF context.



Note

To inter-operate with Cisco routers for [BGP](#) graceful restart, you must enable IPv4 unicast address capability.

Example

The following command disables the route-refresh feature for the peer group outer:

```
disable bgp peer-group outer route-refresh
```

The following command disables the VPNv4 address family for a peer group:

```
disable bgp peer-group backbone capability vpnv4
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Support for Layer 3 VPNs was added in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

disable bgp peer-group next-hop-unchanged

```
disable bgp peer-group peer-group-name l2vpn-evpn next-hop-unchanged
```

Description

Disables a peer group and with respect to the next-hop of routes advertised to EBGp peers.

Syntax Description

<i>peer-group-name</i>	Specifies a peer group.
l2vpn-evpn	Specifies L2VPN EVPN address-family type.
next-hop-unchanged	Enables preserving the BGP next-hop when routes are advertised to EBGp peers (default is disabled).

Default

Default is that next-hop-unchanged is disabled.

Usage Guidelines

This command disables overriding the specification behavior with respect to the next-hop of routes advertised to EBGp peers. Specifically, disabling with this command does not maintain the BGP next-hop for routes advertised to EBGp peers.

Example

The following command disables next-hop unchanged for the BGP peer group pg2 :

```
disable bgp peer-group pg2 l2vpn-evpn next-hop-unchanged
```

History

This command was first available in ExtremeXOS 31.4.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

disable bgp peer-group originate-default

```
disable bgp {peer-group} peer-group-name {address-family [ipv4-unicast |  
ipv4-multicast | ipv6-unicast | ipv6-multicast]} originate-default
```

Description

Removes default routes to all *EBGP* neighbors in the specified peer group.

Syntax Description

<i>peer-group-name</i>	Specifies the BGP peer group for which the default routes are removed.
address-family	Specifies an IPv4 or IPv6 unicast or multicast address family.

Default

Disabled. BGP does not automatically originate and advertise default routes to BGP neighbors.

Usage Guidelines

This command can be successfully executed at any time, irrespective of whether local BGP or the remote BGP peers are enabled or disabled.

If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.



Note

If the specified peer group contains IPv6 peers, it is an IPv6 peer group and you must specify an IPv6 address-family. When the specified peer group is an IPv6 peer group, this command fails if no address family is specified or if an IPv4 address-family is specified. This command also fails if an IPv6 address family is specified for an IPv4 peer-group.

Example

The following command removes default routes for IPv4 unicast traffic for all nodes in the test BGP peer group:

```
disable bgp peer-group test originate-default
```

History

This command was first available in ExtremeXOS 12.2.2.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

disable bgp peer-group remove-private-AS-numbers

```
disable bgp peer-group peer-group-name remove-private-AS-numbers
```

Description

Disables the removal of private autonomous system (AS) numbers from the AS_Path attribute of outbound updates.

Syntax Description

<i>peer-group-name</i>	Specifies a peer group.
------------------------	-------------------------

Default

Disabled.

Usage Guidelines

This command applies to the current VR or VRF context.

Example

The following command disables the *BGP* peer group *outer* from removing private AS numbers:

```
disable bgp peer-group outer remove-private-AS-numbers
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

disable bgp peer-group soft-in-reset

```
disable bgp peer-group peer-group-name {address-family [ipv4-unicast |  
ipv4-multicast | ipv6-unicast | ipv6-multicast] } soft-in-reset
```

Description

Disables the soft input reset feature.

Syntax Description

<i>peer-group-name</i>	Specifies a peer group.
address-family	Specifies an IPv4 or IPv6 unicast or multicast address family.

Default

Disabled.

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

Disabling the soft input reset feature can potentially limit the amount of system memory consumed by the RIB-in.

After you enter this command, the switch automatically disables and enables all neighbors in the peer group before the change takes effect.

This command applies to the current VR or VRF context.

If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.



Note

If the specified peer group contains IPv6 peers, it is an IPv6 peer group and you must specify an IPv6 address-family. When the specified peer group is an IPv6 peer group, this command fails if no address family is specified or if an IPv4 address-family is specified. This command also fails if an IPv6 address family is specified for an IPv4 peer-group.

Example

The following command disables the soft input reset feature:

```
disable bgp peer-group outer soft-in-reset
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the *BGP* feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

disable bgp peer-group

```
disable bgp peer-group peer-group-name
```

Description

Disables a *BGP* peer group and all its BGP neighbors.

Syntax Description

<i>peer-group-name</i>	Specifies a peer group.
------------------------	-------------------------

Default

Disabled.

Usage Guidelines

This command applies to the current VR or VRF context.

Example

The following command disables the BGP peer group outer:

```
disable bgp peer-group outer
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

disable bootp vlan

```
disable bootp { ipv4 } [ dhcp { ipv4 | ipv6 } ] vlan [ vlan | 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

bootp	Disable BOOTP client.
ipv4	IPv4 client. (default)
dhcp	Disable <u>DHCP</u> client.
ipv6	IPv6 client.
vlan	Specify VLAN to configure BOOTP/DHCP client for.
<i>vlan</i>	Specifies a VLAN name.
all	Disables all VLANs.

Default

Disabled.

Usage Guidelines

If the IPv4/IPv6 keyword is not specified, IPv4 is taken as default for the mentioned VLAN.

Example

The following example disables the generation and processing of BOOTP packets on a VLAN named accounting:

```
disable bootp vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

The **ipv4** and **ipv6** keywords were added in ExtremeXOS 15.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable bootprelay ipv6

```
disable bootprelay { ipv4 | ipv6 } {vlan vlan_name} | { vr vr_name} | all
                   {vr vr_name}
```

Description

Disables BOOTP Relay v6. This can be done across the VR or on a per VLAN basis.

Syntax Description

bootprelay	BOOTP Relay service.
IPv4	DHCPv4 BOOTP Relay service.
IPv6	DHCPv6 BOOTP Relay service.
<i>vlan_name</i>	Specifies the VLAN name
<i>vr_name</i>	Specifies the virtual router name.
all	Disables all VLANs.

Default

N/A.

Usage Guidelines

Use this command to disable BOOTP Relay across the VR or on a per VLAN basis.

Example

The following command displays IPv6 bootprelay information:

```
* switch # show bootprelay ipv6
BOOTP Relay: DHCPv6 BOOTP Relay enabled on virtual router "VR-Default"
Include Secondary : Disabled
  BOOTP Relay Servers :2001::1
                      3001::1
                      4001::1
VLAN "Default"      :
  BOOTP Relay       : Enabled
  Interface ID      : 3999 (Default)
  Remote ID         : 00:04:96:52:08:76 (Default)
  Prefix Snooping   : Disabled
VLAN "v1"           :
  BOOTP Relay       : Enabled
  Interface ID      : Interface-Sring1
  Remote ID         :
* switch #
```

History

This command was first available in ExtremeXOS 15.5.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

disable bootprelay

```
disable bootprelay [{vlan} [vlan_name] | [{vr} vr_name] | all [{vr}
vr_name]}
```

Description

Disables the BOOTP relay function on one or all VLANs for the specified VR or VRF.

Syntax Description

<i>vlan_name</i>	Specifies a single VLAN on which to disable the BOOTP relay feature.
<i>vr_name</i>	Specifies a single VR on which to disable the BOOTP relay feature.
all	Specifies that BOOTP relay is to be disabled for all VLANs on the specified VR or VRF.

Default

The BOOTP relay function is disabled on all VLANs and VRs.

Usage Guidelines

Because VLAN names are unique on the switch, you can specify only a VLAN name (and omit the VR name) to disable BOOTP relay. When you disable BOOTP relay on a VR or VRF, BOOTP relay is disabled on all VLANs for that VR. If you enter the command without specifying a VLAN or a VR, the functionality is disabled for all VLANs in the current VR context.

Example

The following command disables the forwarding of BOOTP requests on all VLANs in the current VR context:

```
disable bootprelay
```

You can use either of the following commands to disable the forwarding of BOOTP requests on VLAN unit2:

```
disable bootprelay unit2  
disable bootprelay vlan unit2
```

You can use any one of the following commands to disable the forwarding of BOOTP requests on all VLANs in VR zone3:

```
disable bootprelay zone3  
disable bootprelay vr zone3  
disable bootprelay all zone3  
disable bootprelay all vr zone3
```

History

This command was first available in ExtremeXOS 10.1.

The capability to disable BOOTP relay on a VLAN was added in ExtremeXOS 12.4.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable cdp ports

```
disable cdp ports [port_list | all]
```

Description

Disables CDP on a port.

Syntax Description

<i>port_list</i>	Specifies the list of ports to disable CDP on.
all	Specifies that you disable CDP on all ports.

Default

Enabled.

Usage Guidelines

Example

The following command disables CDP on all ports on the switch:

```
disable cdp ports all
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable cfm segment frame-delay measurement

```
disable cfm segment frame-delay measurement segment_name {mep mep_id}
```

Description

Stops DMM frame transmission.

Syntax Description

<i>segment_name</i>	An alpha numeric string identifying the segment name.
mep <i>mep_id</i>	Specifies the maintenance association End Point that helps trigger a particular MEP level session on that segment. The range is 1-8191. The default is all MEPs on the segment.

Default

N/A.

Usage Guidelines

Use this command to stop transmission of DMM frames for a selected CFM segment. This command stops transmission that has been triggered using the command `enable cfm segment frame-delay measurement`.

This stops the transmission for both continuous and on-demand mode.

Example

The following command stops frame transmission on the CFM segment segment-first:

```
disable cfm frame-delay measurement segment-first
```

History

This command was first available in ExtremeXOS 12.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable cfm segment frame-loss measurement mep

This stops the transmission for both continuous and on-demand mode.

```
disable cfm segment frame-loss measurement segment_name mep mep_id
```

Description

This command stops the transmission of the LMM frames for a particular cfm segment.

Syntax Description

<i>segment_name</i>	An alpha numeric string identifying the segment name.
---------------------	---

Default

N/A.

Usage Guidelines

This below command stops the transmission of the LMM frames for a particular cfm segment. This stops the transmission for both continuous and on-demand mode.

Example

```
disable cfm segment cs2 frame-loss measurement mep 3
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable clear-flow

```
disable clear-flow
```

Description

Disable the CLEAR-Flow agent.

Syntax Description

This command has no arguments or variables.

Default

CLEAR-Flow is disabled by default.

Usage Guidelines

When the CLEAR-Flow agent is disabled, sampling stops and the and all rules are left in the current state. It will not reset actions that were taken while CLEAR-Flow was enabled.

Example

The following example disables CLEAR-Flow on the switch:

```
disable clear-flow
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable cli history expansion

```
disable cli history expansion {session | permanent}
```


Description

Disables command line history expansion.

Syntax Description

cli	Command line interface settings.
history	Command history settings.
expansion	Substitute occurrences of '!\n:w' with the corresponding line 'n' and word 'w+1' from command history (default disabled).
session	Configures history expansion for this CLI session only (default).
permanent	Configures history expansion for this CLI session, and all future sessions.

Default

CLI history expansion is disabled by default.

Usage Guidelines

To view the status of CLI history expansion on the switch, use the `show management` command.

Example

The following command disables CLI history expansion for this session and all future sessions:

```
disable cli history expansion permanent
```

History

This command was first available in ExtremeXOS 32.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable cli prompting

```
disable cli prompting
```

Description

Disables CLI prompting for the session.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

Use this command to have all CLI user prompts automatically continue with the default answer.

This applies to the current session only.

To re-enable CLI prompting for the session, use the `enable cli prompting` command.

To view the status of CLI prompting on the switch, use the `show management` command.

Example

The following command disables prompting:

```
disable cli prompting
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable cli refresh

```
disable cli refresh {session | permanent}
```

Description

This command allows you to disable the default auto refresh behavior. The auto refresh behavior is used for some "show" commands.

Syntax Description

session	Use refresh setting for this CLI session only.
permanent	Use refresh setting for this CLI session, and all future sessions (default).

Default

Permanent.

Usage Guidelines

Use this command to disable the show command auto refresh or add the no-refresh option to the individual command. Since the default for the session may be set to `disable cli refresh`, the commands that take a **no-refresh** option now allow for the alternate **refresh** case if you want to selectively enable a refreshed display.

The **permanent** option is only valid for admin level users.

Example

The following is sample output showing the CLI refresh information.

```
# show management
CLI idle timeout           : Enabled (20 minutes)
CLI max number of login attempts : 3
CLI max number of sessions  : 8
CLI paging                 : Enabled (this session only)
CLI space-completion       : Disabled (this session only)
CLI configuration logging   : Disabled
CLI scripting               : Disabled (this session only)
CLI scripting error mode    : Ignore-Error (this session only)
CLI persistent mode        : Persistent (this session only)
CLI prompting              : Enabled (this session only)
CLI refresh                 : Enabled (this session only)
Telnet access               : Enabled (tcp port 23 vr all)
                           : Access Profile : not set
SSH Access                  : ssh module not loaded.
Web access                  : Enabled (tcp port 80)
                           : Access Profile : not set
```

History

This command was first available in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable cli scripting

```
disable cli scripting {permanent}
```

Description

Disables the use of the CLI scripting commands. When used without the permanent option, it disables the CLI scripting commands for the current session and is a per session setting. The permanent option affects new sessions only and is saved across switch reboots.

Syntax Description

permanent	Disables the CLI scripting commands for new sessions only; this setting is saved across switch reboots.
------------------	---

Default

CLI scripting commands are disabled by default.

Usage Guidelines

You can disable the CLI scripting commands for the session only after this feature has been enabled.

Example

The following command disables the CLI scripting commands for the current session:

```
disable cli scripting
```

History

This command was first available in ExtremeXOS 11.6.

The permanent option was added in ExtremeXOS 12.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable cli scripting output

```
disable cli scripting output
```

Description

Disables the display of CLI commands and responses during script operation.

Syntax Description

This command has no arguments or variables.

Default

During interactive script sessions: CLI scripting output enabled.

During load script command operation: CLI scripting output disabled.

Usage Guidelines

When the CLI scripting output is disabled, the only script output displayed is the `show var {varname}` command and its output. All other commands and responses are not displayed.

When the `load script filename {arg1} {arg2} ... {arg9}` command is entered, the software disables CLI scripting output until the script is complete, and then CLI scripting output is

enabled. Use the enable cli scripting output and disable cli scripting output commands to control what a script displays when you are troubleshooting.

Example

The following command disables CLI scripting output for the current session or until the enable cli scripting output command is entered:

```
disable cli scripting output
```

History

This command was first available in ExtremeXOS 12.1.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable cli space-completion

```
disable cli space-completion
```

Description

Disables the ExtremeXOS feature that completes a command automatically with the spacebar. If you disable this feature, the [Tab] key can still be used for auto-completion.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

None.

Example

The following command disables using the spacebar to automatically complete a command:

```
disable cli space-completion
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

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

Disabled.

Usage Guidelines

Every command is displayed in the log window which allows you to view every command executed on the switch.

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 ExtremeXOS 11.0.

The **cli-config-logging** keyword was split into **cli config-logging** in ExtremeXOS 30.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable cli-config-logging expansion

disable cli-config-logging **expansion**

Description

When CLI logging is enabled, disables showing fully expanded commands, rather than abbreviations, in the log.

Syntax Description

expansion	Disables command expansion in logs.
------------------	-------------------------------------

Default

Expansion is disabled by default.

Usage Guidelines

When CLI logging is enabled (see [enable cli config-logging](#) on page 2095), this command disables showing fully expanded commands, rather than abbreviations, in the log.

For example, with command expansion disabled, a command entered in abbreviated format, such as

```
config por 33 auto of spee 10000 duplex ful
```

appears in the log exactly as it was entered in the command line.

If command expansion is enabled, the command appears in the log in expanded form:

```
configure ports 33 auto off speed 10000 duplex full
```

To see the status of command expansion, use [show management](#) on page 2854.

Example

The following example turns off command expansion:

```
disable cli-config-logging expansion
```

History

This command was first available in ExtremeXOS 22.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable cli paging

```
disable cli paging {session | permanent}
```

Description

Disables pausing at the end of each show screen.

Syntax Description

session	Disables viewing output of commands one screenful at a time for the current user session only (default).
permanent	Disables viewing output of commands one screenful at a time permanently (setting persists after rebooting).

Default

Clipaging is enabled per session by default.

Usage Guidelines

The command line interface (CLI) is designed for use in a VT100 environment.

Most show command output pauses 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.

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 following command disables cli paging permanently (persists after rebooting) and allows you to print continuously to the screen:

```
disable cli paging permanent
```

History

This command was first available in ExtremeXOS 10.1.

The **session** and **permanent** options were added in ExtremeXOS 22.5.

The **clipaging** option was split into two keywords in ExtremeXOS 30.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable cpu-monitoring

```
disable cpu-monitoring
```

Description

Disables CPU monitoring on the switch.

Syntax Description

This command has no arguments or variables.

Default

CPU monitoring is enabled and occurs every 5 seconds.

Usage Guidelines

Use this command to disable CPU monitoring on the switch.

This command does not clear the monitoring interval. Therefore, if you altered the CPU monitoring interval, this command does not return the CPU monitoring interval to 5 seconds. To return to the default frequency level, use the `enable cpu-monitoring {interval seconds} {thresholdpercent}` and specify 5 for the interval.

Example

The following command disables CPU monitoring on the switch:

```
disable cpu-monitoring
```

History

This command was first available in an ExtremeXOS 11.2.

The default value shown began in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable dhcp ports vlan

```
disable dhcp ports port_list vlan vlan_name
```

Description

Disables DHCP on a specified port in a VLAN.

Syntax Description

<i>port_list</i>	Specifies the ports for which DHCP should be disabled.
<i>vlan_name</i>	Specifies the VLAN on whose ports DHCP should be disabled.

Default

N/A.

Usage Guidelines

None.

Example

The following command disables DHCP for port 6:9 in VLAN corp:

```
disable dhcp ports 6:9 vlan corp
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable dhcp vlan

```
disable dhcp [ipv4 | ipv6] vlan [vlan_name | all]
```

Description

Disables the generation and processing of *DHCP* packets on a *VLAN* to obtain an IP address for the VLAN from a DHCP server.

Syntax Description

<i>vlan_name</i>	Specifies a VLAN name.
all	Specifies all VLANs.

Default

If the IPv4/IPv6 keyword is not specified, IPv4 is taken as default for the mentioned VLAN.

Usage Guidelines

None.

Example

The following command disables the generation and processing of DHCP packets on a VLAN named accounting:

```
disable dhcp vlan accounting
disable dhcp ipv6 vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

This command was modified in ExtremeXOS 15.6 to include the **ipv4** and **ipv6** keywords

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable diffserv examination ports

```
disable diffserv examination ports [port_list | all]
```

Description

Disables the examination of the DiffServ field in an IP packet.

Syntax Description

<i>port_list</i>	Specifies a list of ports or slots and ports to which the parameters apply.
all	Specifies that DiffServ examination should be disabled for all ports.

Default

Disabled.

Usage Guidelines

The diffserv examination feature is disabled by default.

Example

The following command disables DiffServ examination on the specified ports:

```
disable diffserv examination ports 5:3,5:5,6:6
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable diffserv replacement ports

```
disable diffserv replacement ports [port_list | all] [{qosprofile}  
                                qosprofile}
```

Description

Disables the replacement of DiffServ code points in packets transmitted by the switch.

Syntax Description

<i>port_list</i>	Specifies a list of ports or slots and ports on which Diffserv replacement will be disabled.
all	Specifies that DiffServ replacement should be disabled for all ports.
qosprofile	Disables DiffServ on a QoS profile. Note: If this option is not specified it will disable DiffServ replacement on all qosprofiles.
<i>qosprofile</i>	Specifies the QoS profile number.

Default

The DiffServ replacement feature is disabled by default.

Usage Guidelines

N/A.

Example

The following example disables DiffServ replacement on selected ports:

```
disable diffserv replacement ports 1:2,5:5,6:6
```

History

This command was first available in ExtremeXOS 11.0.

The **gosprofile** keyword and *gosprofile* variable were added in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable dns cache

```
disable dns cache [{vlan} vlan_name | {vr} vr_name]
```

Description

Disables the Domain Name System (DNS) cache on a virtual router (VR) or VLAN.

Syntax Description

dns	Domain name system.
cache	Specifies disabling the DNS cache.
vlan	Specifies disabling DNS cache on a VLAN.
<i>vlan_name</i>	Specifies the VLAN name.
vr	Specifies disabling DNS cache on a VR.
<i>vr_name</i>	Specifies the VR name. If not specified, the VR of the current command context is used.

Default

If no VR name is specified, the VR of the current command context is used.

Usage Guidelines

To view the DNS cache configuration, use the command `show dns cache configuration [{vlan} vlan_name | {vr} vr_name]`

Example

The following example disables DNS cache on VLAN "VLAN1":

```
# disable dns cache vlan VLAN1
```

History

This command was first available in ExtremeXOS 30.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable dns cache analytics

```
disable dns cache analytics [{vr} vr_name]
```

Description

Disables Domain Name System (DNS) analytics.

Syntax Description

dns	Domain Name System.
cache	Specifies DNS cache.
analytics	Specifies disabling DNS cache analytics. Analytics provides more insight into DNS queries when DNS cache is enabled. Default is disabled.
vr	Specifies disabling DNS analytics on a VR.
<i>vr_name</i>	Specifies the VR name. If not specified, the VR of the current command context is used.

Default

DNS analytics is disabled by default.

Usage Guidelines

To enable DNS analytics, use the command `enable dns cache analytics [{vr} vr_name]`.

Example

The following example disables DNS analytics on VR "vr1":

```
# disables dns cache analytics vr vr1
```

History

This command was first available in ExtremeXOS 30.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable dns cache dnssec

```
disable dns cache {dnssec}
```

Description

Disables validating DNS replies and cache data for DNSSEC (Domain Name System Security Extensions).

Syntax Description

dnssec	Disables validating DNS replies and cache data for DNSSEC. Default is disabled.
---------------	---

Default

By default, DNSSEC is disabled.

Usage Guidelines

You cannot disable DNSSEC if DNS cache is enabled.

Example

The following example disables DNSSEC:

```
# disable dns cache dnssec
```

History

This command was first available in ExtremeXOS 31.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable dos-protect

```
disable 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 dos-protect
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable dot1p examination inner-tag ports

disable dot1p examination inner-tag ports [**all** | *port_list*]

Description

Used with VMANs, and instructs the switch to examine the 802.1p value of the outer tag, or added VMAN header, to determine the correct egress queue on the egress port.

Syntax Description

all	Specifies all ports.
<i>port_list</i>	Specifies a list of ports or slots and ports.

Default

Disabled.

Usage Guidelines

Use this command to instruct the system to refer to the 802.1p value contained in the outer tag, or VMAN encapsulation tag, when assigning the packet to an egress queue at the egress port of the VMAN.



Note

See “Quality of Service” in the [Switch Engine 32.3 User Guide](#) for information on configuring and displaying the current 802.1p and DiffServ configuration for the inner, or original header, 802.1p value.

Example

The following example uses the 802.1p value on the outer tag, or VMAN encapsulation, to put the packet in the egress queue on the VMAN egress port:

```
disable dot1p examination inner-tag port 3:2
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable dot1p examination ports

```
disable dot1p examination ports [port_list | all]
```

Description

Prevents examination of the 802.1p priority field as part of the [QoS](#) configuration.

Syntax Description

<i>port_list</i>	Specifies a list of ports or slots and ports.
all	Specifies that dot1p replacement should be disabled for all ports.

Default

Enabled.

Usage Guidelines

The 802.1p examination feature is enabled by default. To free *ACL* resources, disable this feature whenever another QoS traffic grouping is configured. (For information on available ACL resources, see *ACLs* in the [Switch Engine 32.3 User Guide](#))



Note

If you disable this feature when no other QoS traffic grouping is in effect, 802.1p priority enforcement of 802.1q tagged packets continues.

SummitStack Only.

Dot1p examination cannot be disabled for priority values 5 and 6. However, the precedence of the examination is lowered so that all other traffic grouping precedences are higher. The mappings you configure with the `configure dot1p type` command remain in effect.

As part of the COS global status enable action, COS will automatically enable dot1p examination on all ports. An internal status will track this event. The `disable dot1p examination` command will print an additional warning message in the event that COS was configured via *SNMP*. If the COS global status is disabled via SNMP, the internal status will be cleared and the additional WARNING message will not be displayed.

Example

The following command disables 802.1p value examination on ports 1 to 5:

```
disable dot1p examination ports 1-5
```

History

This command was available in ExtremeXOS 11.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable dot1p replacement ports

```
disable dot1p replacement ports [port_list | all] [{qosprofile}  
  qosprofile}
```

Description

Disables the ability to overwrite 802.1p priority values for a given set of ports.

Syntax Description

<i>port_list</i>	Specifies a list of ports or slots and ports to which the parameters apply.
all	Specifies that 802.1p replacement should be disabled for all ports.
qosprofile	Disables 802.1p on a QoS profile. Note: If this option is not specified it will disable dot1p replacement for all qosprofiles.
<i>qosprofile</i>	Specifies the QoS profile number.

Default

N/A.

Usage Guidelines

The dot1p replacement feature is disabled by default.

Beginning with ExtremeXOS version 11.4 on the 1 Gigabit Ethernet ports, 802.1p replacement always happens when you configure the DiffServ traffic grouping.

Example

The following example disables 802.1p value replacement on all ports:

```
disable dot1p replacement ports all
```

History

This command was first available in ExtremeXOS 11.0.

The **qosprofile** keyword and *qosprofile* variable were added in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable eaps

```
disable eaps {name}
```

Description

Disables the [EAPS](#) function for a named domain or for an entire switch.

Syntax Description

<i>name</i>	Specifies the name of an EAPS domain.
-------------	---------------------------------------

Default

Disabled for the entire switch.

Usage Guidelines

To prevent loops in the network, the switch displays by default a warning message and prompts you to disable EAPS for a specific domain or the entire switch. When prompted, do one of the following:

- Enter y to disable EAPS for a specific domain or the entire switch.
- Enter n or press [Return] to cancel this action.

If you have considerable knowledge and experience with EAPS, you might find the EAPS loop protection warning messages unnecessary. For more information, see the [configure eaps config-warnings off](#).

Example

The following command disables the EAPS function for entire switch:

```
disable eaps
```

The switch displays the following warning message and prompts you to confirm this action:

```
WARNING: Disabling EAPS on the switch could cause a loop in the network!  
Are you sure you want to disable EAPS? (y/n) Enter y to disable EAPS on the switch. Enter n to cancel  
this action.
```

The following command disables the EAPS function for the domain eaps-1:

```
disable eaps eaps-1
```

The switch displays the following warning message and prompts you to confirm this action:

```
WARNING: Disabling specific EAPS domain could cause a loop in the  
network!
```

```
Are you sure you want to disable this specific EAPS domain? (y/n)
```

Enter y to disable the EAPS function for the specified domain. Enter n to cancel this action.

History

This command was first available in ExtremeXOS 11.0.

The interactive messages were added in ExtremeXOS 11.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable edp ports

```
disable edp ports [ports | all]
```

Description

Disables the *EDP* on one or more ports.

Syntax Description

<i>ports</i>	Specifies one or more ports or slots and ports, including management port.
all	Specifies all ports on the switch, including management port.

Default

Enabled.

Usage Guidelines

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.

Example

The following command disables EDP on ports 2 and 4 on a switch:

```
disable edp ports 2,4
```

History

This command was first available in ExtremeXOS 10.1.

Ability to disable EDP on management port was added in ExtremeXOS 30.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable elrp-client

```
disable elrp-client
```

Description

Disables the ELRP client (standalone ELRP) globally.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command disables the ELRP globally so that none of the ELRP [VLAN](#) configurations take effect.

The ELRP client must be enabled globally in order for it to work on any VLANs. Use the [enable elrp-client](#) command to globally enable the ELRP client.

Example

The following command globally disables the ELRP client:

```
disable elrp-client
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable elsm ports

```
disable elsm ports port_list
```

Description

Disables the ELSM protocol for the specified ports.

Syntax Description

<i>port_list</i>	Specifies the port or ports for which ELSM should be disabled.
------------------	--

Default

The default is disabled.

Usage Guidelines

ELSM works between two connected ports, and each ELSM instance is based on a single port. When you disable ELSM on the specified ports, the ports no longer send ELSM hello messages to their peers and no longer maintain ELSM states.

When you enable ELSM on the specified ports, the ports participate in ELSM with their peers and begin exchanging ELSM hello messages. To enable ELSM, use the following command:

```
enable elsm ports port_list
```

For more information about ELSM, see the command `enable elsm ports`.

Example

The following command disables ELSM for slot 2, ports 1-2 on the switch:

```
disable elsm ports 2:1-2:2
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable elsm ports auto-restart

```
disable elsm ports port_list auto-restart
```

Description

Disable ELSM automatic restart for the specified ports.

Syntax Description

<i>port_list</i>	Specifies the port or ports for which ELSM auto-restart is being disabled.
------------------	--

Default

The default is enabled.

Usage Guidelines

If you disable ELSM automatic restart, the ELSM-enabled port can transition between the following states multiple times: Up, Down, and Down-Wait. When the number of state transitions is greater than or equal to the sticky threshold, the port enters and remains in the Down-Stuck state.

The ELSM sticky threshold specifies the number of times a port can transition between the Up and Down states. The sticky threshold is not user-configurable and has a default value of 1. That means a port can transition only one time from the Up state to the Down state. If the port attempts a subsequent transition from the Up state to the Down state, the port enters the Down-Stuck state.

If the port enters the Down-Stuck state, you can clear the stuck state and have the port enter the Down state by using one of the following commands:

```
clear elsm ports port_list auto-restart
```

```
enable elsm ports port_list auto-restart
```

If you use the `enable elsm ports` command, automatic restart is always enabled; you do not have to use the `clear elsm ports` command to clear the stuck state.

Enabling Automatic Restart

To enable ELSM automatic restart, you must explicitly configure this behavior on each ELSM-enabled port. If you enable ELSM automatic restart and an ELSM-enabled port goes down, ELSM bypasses the Down-Stuck state and automatically transitions the down port to the Down state, regardless of the number of times the port goes up and down.

To enable automatic restart, use the following command:

```
enable elsm ports port_list auto-restart
```

If you configure automatic restart on one port, we recommend that you use the same configuration on its peer port.

Example

The following example disables ELSM automatic restart for slot 2, ports 1-2 on the switch:

```
disable elsm ports 2:1-2:2 auto-restart
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable erps

```
disable erps
```

Description

Disable *ERPS* (ITU-T G.8032 standard).

Syntax Description

N/A.

Default

N/A.

Usage Guidelines

Use this command to disable ERPS.

Example

The following command disables ERPS:

```
disable erps
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

disable erps block-vc-recovery

disable erps *ring-name* **block-vc-recovery**

Description

Disables the ability on *ERPS* rings to block virtual channel recovery to avoid temporary loops .

Syntax Description

<i>ring-name</i>	Alphanumeric string that identifies the ERPS ring.
block-vc-recovery	Block on Virtual channel recovery.

Default

N/A.

Usage Guidelines

Use this command to disable the ability on ERPS rings to block on virtual channel recovery to avoid temporary loops. This is done on interconnected nodes for sub-ring configurations.

Example

The following example disables a virtual channel recovery block on “ring1”:

```
diabile erps ring1 block-vc-recovery
```

History

This command was first available in ExtremeXOS 15.13.

Platform Availability

This command is available on all platforms that are running ExtremeXOS.

disable erps ring-name

```
disable erps ring-name
```

Description

Disable an existing *ERPS* ring/sub-ring.

Syntax Description

<i>ring-name</i>	Alphanumeric string that identifies the ERPS ring.
------------------	--

Default

N/A.

Usage Guidelines

Use this command to disable an existing ERPS ring/sub-ring.

Example

The following example disables an existing ERPS ring identified as “ring1”:

```
disable erps ring1
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

disable erps topology-change

disable erps *ring-name* **topology-change**

Description

Disable the ability of *ERPS* to set the topology-change bit to send out Flush events.

Syntax Description

<i>ring-name</i>	Alphanumeric string that identifies the ERPS sub-ring.
topology-change	Topology change propagation control.

Default

N/A.

Usage Guidelines

Use this command to disable the ability of ERPS to set the topology-change bit to send out Flush events.

Example

The following example disables the ability to set the topology-change bit for an existing ERPS sub-ring identified as “ring1”:

```
disable erps ring1 topology-change
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

disable esrp

disable esrp {*esrpDomain*}

Description

Disables *ESRP* for a named domain or for the entire switch.

Syntax Description

<i>esrpDomain</i>	Specifies the name of an ESRP domain.
-------------------	---------------------------------------

Default

Disabled for the entire switch.

Usage Guidelines

If you do not specify a domain name, ESRP is disabled for the entire switch.

If you disable an ESRP domain, the domain enters the Aware state, the switch notifies its neighbor that the ESRP domain is going down, and the neighbor clears its neighbor table. If the master switch receives this information, it enters the neutral state to prevent a network loop. If the slave switch receives this information, it enters the neutral state.

Example

The following command disables ESRP for the entire switch:

```
disable esrp
```

The following command disables ESRP for the domain esrp1:

```
disable esrp esrp1
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms that use the Edge, Advanced Edge, or Core license. For information on the licenses available for each platform, see the [Switch Engine 32.3 Feature License Requirements](#) document.

disable ethernet oam ports link-fault-management

```
disable ethernet oam ports [port_list | all] link-fault-management
```

Description

Disables Ethernet OAM on ports.

Syntax Description

<i>port_list</i>	Specifies the particular ports.
all	Specifies all fiber ports.

Default

Ethernet OAM is disabled on all ports.

Usage Guidelines

Use this command to disable Ethernet OAM on one or more specified ports or on all fiber ports.

When operating as a stack master, the ExtremeSwitching switch can process this command for ports on supported platforms.

Example

The following command disables Ethernet OAM on port 1:

```
# disable ethernet oam ports 1 link-fault-management
```

History

This command was first available in ExtremeXOS 12.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable fdb static-mac-move

```
disable fdb static-mac-move
```

Description

Disables EMS and *SNMP* reporting of discovered MAC addresses that are duplicates of statically configured MAC addresses.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

None.

Example

The following example disables this feature:

```
disable fdb static-mac-move
```

History

This command was first available in ExtremeXOS 12.7.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable flooding ports

With this command you can further identify the type of packets for which to block flooding.

```
disable flooding [all_cast | broadcast | multicast | unicast] ports  
[port_list | all]
```

Description

Disables Layer 2 egress flooding on one or more ports.

Syntax Description

all_cast	Specifies disabling egress flooding for all packets on specified ports.
broadcast	Specifies disabling egress flooding only for broadcast packets.
multicast	Specifies disabling egress flooding only for multicast packets.
unicast	Specifies disabling egress flooding only for unknown unicast packets.
<i>port_list</i>	Specifies one or more ports or slots and ports.
all	Specifies all ports on the switch.

Default

Enabled for all packet types.

Usage Guidelines



Note

If an application requests specific packets on a specific port, those packets are not affected by the `disable flooding ports` command.

You might want to disable egress flooding to do the following:

- enhance security
- enhance privacy

- improve network performance

This is particularly useful when you are working on an edge device in the network. The practice of limiting flooded egress packets to selected interfaces is also known as upstream forwarding.

**Note**

If you disable egress flooding with static MAC addresses, this can affect many protocols, such as IP and ARP.

The following guidelines apply to enabling and disabling egress flooding:

- Disabling multicasting egress flooding does not affect those packets within an *IGMP* membership group at all; those packets are still forwarded out. If IGMP snooping is disabled, multicast packets are not flooded.
- Egress flooding can be disabled on ports that are in a load-sharing group. In a load-sharing group, the ports in the group take on the egress flooding state of the master port; each member port of the load-sharing group has the same state as the master port.
- On all platforms *FDB* learning takes place on ingress ports and is independent of egress flooding; either can be enabled or disabled independently.
- Disabling unicast or all egress flooding to a port also stops packets with unknown MAC addresses to be flooded to that port.
- Disabling broadcast or all egress flooding to a port also stops broadcast packets to be flooded to that port.

You can disable egress flooding for unicast, multicast, or broadcast MAC addresses, as well as for all packets on the ports of the switch. The default behavior is enabled egress flooding for all packet types.

Example

The following example disables unicast flooding on ports 10-12::

```
# disable flooding unicast port 10-27
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable flow-control ports

```
disable flow-control [tx-pause {priority priority} | rx-pause  
                    {qosprofile qosprofile}] ports [all | port_list]
```

Description

Disables specified flow control configurations.

Syntax Description

tx-pause	Specifies transmission pause processing.
<i>priority</i>	Specifies all priorities or single priorities--dot1p priority for tagged packets and internal priority for untagged packets. Used with priority flow control only.
rx-pause	Specifies reception pause processing.
<i>qosprofile</i>	Specifies a QoS profile ("qp1" "qp2" "qp3" "qp4" "qp5" "qp6" "qp7" "qp8") to pause for priority flow control packet reception. Used with priority flow control only.
all	Specifies all ports or slots.
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

Disabled.

Usage Guidelines

IEEE 802.3x-Flow Control

Use this command to disable the processing of IEEE 802.3x pause flow control messages received from the remote partner. Disabling rx-pause processing avoids dropping packets in the switch and allows for better overall network performance in some scenarios where protocols such as TCP handle the retransmission of dropped packets by the remote partner.

To disable RX flow-control, TX flow-control must first be disabled. Refer to the [disable flow-control ports](#) command. If you attempt to disable RX flow-control with TX flow-control enabled, an error message is displayed.

IEEE 802.1Qbb-Priority Flow Control

Use this command to disable the processing of IEEE 802.1Qbb priority flow control messages received from the remote partner. Disabling TX stops the port from transmitting PFC packets for that priority, regardless of congestion. Disabling RX stops the processing of PFC packets received on that port for the specific QoS profile.

Example

IEEE 802.3x

The following command disables the tx flow-control feature on ports 5 through 7 on an ExtremeSwitching switch:

```
# disable flow-control tx-pause ports 5-7
```

IEEE 802.1Qbb

The following command disables TX for priority 3 on port 3:

```
# disable flow-control tx-pause priority 3 ports 3
```

The following command disables RX for QoS profile qp4 on port 6:

```
# disable flow-control rx-pause qosprofile qp4 port 6
```

History

This command was first available in ExtremeXOS 12.1.3.

The priority function (PFC) was added in ExtremeXOS 12.5.

Platform Availability

IEEE 802.3x

The basic TX-pause and RX-pause functions of this command are available on all switches.

IEEE 802.1Qbb

The priority function (PFC) is available only on 10G ports.

disable flowmon

```
disable flowmon
```

Description

Disables Flow Monitor.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

Modification rules for groups, keys, and collectors remain active after Flow Monitor is disabled. A new group can be created and configured with its parameters, collector, template portions, and added keys. The group can also be enabled or disabled.

Enabling a group while Flow Monitor is disabled will program the hardware, but the flow collection for the group will be disabled.

Example

The following command disables Flow Monitor:

```
# disable flowmon
```

History

This command was first available in ExtremeXOS 32.2.

Platform Availability

This command is available on ExtremeSwitching 5520 and 5420 series switches.

disable flowmon group

```
disable flowmon group group_name
```

Description

Disables a Flow Monitor group.

Syntax Description

group	Specifies the Flow Monitor group.
<i>group_name</i>	Specifies the assigned name of the Flow Monitor group. Range is 32 characters.

Default

N/A.

Usage Guidelines

Example

The following command disables a Flow Monitor group with the name 'max-flow-age':

```
# disable flowmon group max-flow-age
```

History

This command was first available in ExtremeXOS 32.2.

Platform Availability

This command is available on ExtremeSwitching 5520 and 5420 series switches.

disable icmp ipv6

```
disable icmp ipv6 [ignore-multicasts | ignore-anycasts]
```

Description

Disables the *ICMP* IPv6 reply to multicast or anycast echo request.

Syntax Description

ignore-multicasts	Specifies to reply to ICMP echo requests destined to an IP multicast address. Default is ignore (disable).
ignore-anycasts	Specifies to reply to ICMP echo requests destined to an IP anycast address. Default is ignore (disable).

Default

Ignore (disable).

Usage Guidelines

Use this command to disable ignoring a reply packet to multicast or anycast echo request.

Example

The following example specifies to reply to ICMP multicast echo requests:

```
disable icmp ipv6 ignore-multicasts
```

History

This command was first available in ExtremeXOS 31.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable icmp redirects ipv6 fast-path

```
disable icmp redirects ipv6 fast-path
```

Description

When disabled (default), only slow path packets (packets that cannot be forwarded by hardware) may trigger *ICMP* redirects.

Syntax Description

fast-path	Only slow path packets (packets that cannot be forwarded by hardware) may trigger ICMP redirects.
------------------	---

Default

Disabled.

Usage Guidelines

Use this command so that only slow path packets (packets that cannot be forwarded by hardware) may trigger ICMP redirects.

Example

The enabled or disabled setting is displayed when entering the command:

```
# show ipconfig ipv6
Route Sharing           : Disabled
ICMP Redirect for Fast Path : Enabled
Max Shared Gateways     : Current: 4   Configured: 4

Interface      IPv6 Prefix      Flags
v1             2001::1/24             -EUf---R-
v1             fe80::204:96ff:fele:ec00%v1/64 -EUfP--R-
Flags : D - Duplicate address detected on VLAN, T - Tentative address
E - Interface enabled, U - Interface up, f - IPv6 forwarding enabled,
i - Accept received router advertisements enabled,
R - Send redirects enabled, r - Accept redirects enabled
P - Prefix address
BD-8810.2 #
```

History

This command was first available in ExtremeXOS 15.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable icmp redirects

```
disable icmp redirects {ipv4} {vlan all | {vlan} {name}}
```

Description

Disables the generation of *ICMP* redirect messages on one or all *VLANs*.

Syntax Description

<i>name</i>	Specifies a VLAN name.
-------------	------------------------

Default

Enabled.

Usage Guidelines

Disables the generation of ICMP redirects (type 5) to hosts who direct routed traffic to the switch where the switch detects that there is another router in the same subnet with a better route to the destination.

Example

The following example disables ICMP redirects from VLAN "accounting":

```
disable icmp redirects vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable icmp userredirects

```
disable icmp userredirects
```

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.

If the switch has a route to a destination network, the switch uses that router as the gateway to forward the packets to. If that router knows about a better route to the destination, and the next hop is in the same subnet as the originating router, the second router sends an ICMP redirect message to the first router. If ICMP useredirects is disabled, the switch disregards these messages and continues to send the packets to the second router.

Example

The following example disables the changing of routing table information:

```
disable icmp useredirects
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable identity-management

```
disable identity-management
```

Description

Disables the identity management feature, which tracks users and devices that connect to the switch.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

Only admin-level users can execute this command.



Note

If the identity management feature is running and then disabled, all identity management database entries are removed and cannot be retrieved. If identity management is enabled later, the identity management feature starts collecting information about currently connected users and devices.

Example

The following command disables the identity management feature:

```
disable identity-management
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable cli idletimeout

```
disable cli idle-timeout
```

Description

Disables the timer that disconnects idle sessions from the switch.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Timeout 20 minutes.

Usage Guidelines

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

Telnet sessions remain open until you close the Telnet client.

If you have an SSH2 session and disable the idle timer, the SSH2 connection times out after 61 minutes of inactivity.

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 cli idle-timeout
```

History

This command was first available in ExtremeXOS 10.1.

The **cli** keyword was added and the **idletimeout** keyword was changed to **idle-timeout** in ExtremeXOS 30.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable igmp

```
disable igmp {vlan name}
```

Description

Disables IGMP on a router interface. If no VLAN is specified, IGMP is disabled on all router interfaces.

Syntax Description

<i>name</i>	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, hosts respond 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 and processing of IGMP packets. IGMP should be enabled when the switch is configured to perform IP multicast routing.

This command disables IGMPv2 and IGMPv3.

Example

The following example disables IGMP on VLAN accounting:

```
disable igmp vlan accounting
```


History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IPv4 multicast feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

disable igmp snooping vlan fast-leave

```
disable igmp snooping {vlan} name fast-leave
```

Description

Disables the [*IGMP*](#) snooping fast leave feature on the specified [*VLAN*](#).

Syntax Description

<i>name</i>	Specifies a VLAN.
-------------	-------------------

Default

Disabled.

Usage Guidelines

None.

Example

The following command disables the IGMP snooping fast leave feature on the default VLAN:

```
disable igmp snooping "Default" fast-leave
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IGMP snooping feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

disable igmp snooping

```
disable igmp snooping {forward-mcrouter-only | with-proxy | vlan name}
```

Description

Disables IGMP snooping.

Syntax Description

forward-mcrouter-only	Specifies that the switch forwards all multicast traffic to the multicast router only.
with-proxy	Disables the IGMP snooping proxy.
<i>name</i>	Specifies a <u>VLAN</u> .

Default

IGMP snooping and the with-proxy option are enabled by default, but forward-mcrouter-only option is disabled by default.

Usage Guidelines

If a VLAN is specified, IGMP snooping is disabled only on that VLAN, otherwise IGMP snooping is disabled on all VLANs.

This command applies to both IGMPv2 and IGMPv3.

If the switch is in the forward-mcrouter-only mode, then the command `disable igmp snooping forward-mcrouter-only` changes the mode so that all multicast traffic is forwarded to any IP router. If not in the forward-mcrouter-mode, the command `disable igmp snooping forward-mcrouter-only` has no effect.

To change the snooping mode you must disable IP multicast forwarding. Use the command: `disable ipmcforwarding`

The **with-proxy** option can be used for troubleshooting purpose. It should be enabled for normal network operation.

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 last member leaves the group.

Example

The following example disables IGMP snooping on the VLAN accounting:

```
disable igmp snooping accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IGMP snooping feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

disable igmp ssm-map

```
disable igmp ssm-map {vr vr-name}
```

Description

Disables IGMP SSM mapping.

Syntax Description

<i>vr-name</i>	Specifies a virtual router name. If the VR name is omitted, the switch disables mapping on the VR specified by the current CLI VR context.
----------------	--

Default

Disabled on all interfaces.

Usage Guidelines

None.

Example

The following command disables IGMP-SSM mapping on the VR in the current CLI VR context:

```
disable igmp ssm-map
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IPv4 multicast feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

disable inline-power

```
disable inline-power [{fast {ports [port_list | all]}} | perpetual]
```

Description

Disables PoE, and perpetual PoE to all ports; or fast PoE to all ports, or selected ports, for some platforms.

Syntax Description

fast	Disables delivery of PoE power to devices at the time of switch power on without waiting for boot up based on last saved PoE state. The default is disabled.
ports	For fast PoE, specifies selecting ports. ExtremeSwitching 5320, 5420, 5520, and 5720 series switches only.
<i>port_list</i>	For fast PoE, specifies the port list separated by a comma or -. ExtremeSwitching 5320, 5420, 5520, and 5720 series switches only.
all	For fast PoE, specifies selecting all ports. ExtremeSwitching 5320, 5420, 5520, and 5720 series switches only.
perpetual	Disable preserving PoE power delivery to devices during reboot. Perpetual PoE is a switch-wide setting. The default is disabled.

Default

By default:

- PoE is enabled.
- Fast PoE is disabled.
- Perpetual PoE is disabled.

Usage Guidelines

You can control whether inline power is provided to the system by using the `disable inline-power` command and the `enable inline power` command. Using the `disable inline-power` command shuts down inline power currently provided on the entire switch or to specified ports and slots. Disabling inline power to a switch, port, or slot immediately removes power to any connected PDs. By default, inline power provided to all ports is enabled. Additionally, you can disable delivery of PoE power to devices at the time of switch power on without waiting for boot up (fast PoE) based on last saved PoE state. Per-port fast PoE is available on certain platforms. You can also elect to not preserve PoE power delivery to devices during reboot (perpetual PoE). The default for both PoE options is disabled.



Note

Disabling inline power using the `disable inline-power` command does not affect the data traffic traversing the port. And, disabling the port using the `disable port` command does not affect the inline power supplied to the port.

**Note**

Inline power cannot be delivered to connected PDs unless the switch is powered on.

Example

The following command shuts down inline power currently provided to all ports and all slots:

```
disable inline-power
```

The following example turns off perpetual PoE for the switch:

```
# disable inline-power perpetual
```

The following example turns off fast PoE for ports 1,2, and 5:

```
# disable inline-power fast ports 1,2,5
```

History

This command was first available in ExtremeXOS 11.1.

The **fast** and **perpetual** PoE options were added in ExtremeXOS 30.3.

Per-port fast PoE was added for ExtremeXOS 31.1.

Platform Availability

This command is available on the PoE devices listed in *Extreme Networks PoE Devices* in the [Switch Engine 32.3 User Guide](#).

The **fast** and **perpetual** options are only available on the ExtremeSwitching 5320, 5420, 5520, and 5720 (per port) series switches.

disable inline-power ports

```
disable inline-power ports [all | port_list]
```

Description

Shuts down *PoE* power currently provided to all ports or to specified ports.

Syntax Description

all	Disables inline power to all ports on the switch.
<i>port_list</i>	Disables inline power to the specified ports.

Default

Enable.

Usage Guidelines

Disabling inline power to ports immediately removes power to any connected PDs. By default, the capability to provide inline power to all ports is enabled.



Note

Disabling inline power using the `disable inline-power` command does not affect the data traffic traversing the port. And, disabling the port using the `disable port` command does not affect the inline power supplied to the port.

Disabling inline power to a port providing power to a PD immediately removes power to the PD.

Example

The following command shuts down inline power currently provided to ports 4 and 5 on a switch:

```
disable inline-power ports 4,5
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on the PoE devices listed in [Extreme Networks PoE Devices](#).

disable inline-power slot

```
disable inline-power [fast | perpetual] slot slot
```

Description

Shuts down *PoE*, and fast and perpetual PoE, power currently provided to the specified slot.

Syntax Description

<i>slot</i>	Selects the slot to disable inline power, or fast/perpetual PoE power on.
fast	Disable delivery of PoE power to devices at the time of switch power on without waiting for boot up based on last saved PoE state. The default is disabled.
perpetual	Disable preserving PoE power delivery to devices during reboot. Perpetual PoE is a switch-wide setting. The default is disabled.

Default

By default:

- PoE is enabled.

- Fast PoE is disabled.
- Perpetual PoE is disabled.

Usage Guidelines

Disabling inline power to a slot immediately removes power to any connected PDs. By default, the capability to provide inline power to a slot is enabled. Additionally, you can disable delivery of PoE power to devices at the time of switch power on without waiting for boot up (fast PoE) based on last saved PoE state. You can also elect to not preserve PoE power delivery to devices during reboot (perpetual PoE). The default for both PoE options is disabled.



Note

You can set the reserved power budget to 0 for a slot if, and only if, you first issue this command.

On a stack if you do not specify a slot number, the command operates on all active nodes. This command operates only on nodes in the active topology.

Example

The following command removes power to all PDs on slot 3:

```
disable inline-power slot 3
```

The following example turns off perpetual PoE for slot 3:

```
# disable inline-power perpetual slot 3
```

History

This command was first available in ExtremeXOS 11.1.

The **fast** and **perpetual** PoE options were added in ExtremeXOS 30.3.

Platform Availability

This command is available on the PoE devices listed in [Extreme Networks PoE Devices](#).

disable ip anycast

```
disable ip anycast {vlan} vlan_name
```

Description

Disables IP anycast on a VLAN.

Syntax Description

ip	Layer 3 Internet Protocol.
anycast	Disables IP anycast on a VLAN.
vlan	Selects the VLAN.
<i>vlan_name</i>	Specifies the VLAN name.

Default

N/A.

Usage Guidelines

N/A.

Example

The following example disables IP anycast on the VLAN "vlan1":

```
# disable ip anycast vlan vlan1
```

History

This command was first available in ExtremeXOS 30.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable ip nat

```
disable ip nat
```

Description

Globally disables Network Address Translation (NAT).

Syntax Description

ip	Specifies Internet Protocol (IP).
nat	Specifies enabling NAT.

Default

N/A.

Usage Guidelines

NAT maps IP addresses from one address domain (typically private IP address spaces) to an another address domain (typically a public Internet IP address space) to provide transparent routing to end hosts. This translation is accomplished transparently by having a NAT device translate the IP address and/or Layer 4 port of the packets.

To view IP NAT information, run the command `show ip nat`.

Example

The following example disables IP NAT:

```
# disable ip nat
```

History

This command was first available in ExtremeXOS 31.2.

Platform Availability

This command is available on the switches that support the IP NAT feature. For information about which switches support this and other features, see the [Switch Engine 32.3 Feature License Requirements](#) document.

disable ip nat rule

```
disable ip nat rule rule_name
```

Description

Disables Network Address Translation (NAT) rules.

Syntax Description

ip	Specifies Internet Protocol (IP).
nat	Specifies NAT.
rule	Specifies disabling a NAT rule.
<i>rule_name</i>	Specifies the NAT rule to disable.

Default

N/A.

Usage Guidelines

You must disable a rule to make configuration changes to it.

Example

The following example the disables the IP NAT rule "rule1":

```
# disables ip nat rule rule1
```

History

This command was first available in ExtremeXOS 31.2.

Platform Availability

This command is available on the switches that support the IP NAT feature. For information about which switches support this and other features, see the [Switch Engine 32.3 Feature License Requirements](#) document.

disable iparp checking

```
disable iparp {vr vr_name} checking
```

Description

Disable checking if the ARP request source IP address is within the range of the local interface or VLAN domain.

Syntax Description

<i>vr_name</i>	Specifies a VR or VRF.
----------------	------------------------

Default

Enabled.

Usage Guidelines

If you do not specify a VR or VRF, the command applies to the current VR context.

Example

The following example disables IP ARP checking:

```
disable iparp checking
```

History

This command was first available in ExtremeXOS 10.1.

The **vr** option was added in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable iparp gratuitous protect vlan

```
disable iparp gratuitous protect [ {vlan} vlan_name | vlan vlan_list]
```

Description

Disables gratuitous ARP protection on the specified VLAN.

Syntax Description

<i>vlan_name</i>	Specifies the VLAN.
<i>vlan_list</i>	Specifies a VLAN list of IDs.

Default

Disabled.

Usage Guidelines

Hosts can launch man-in-the-middle attacks by sending out gratuitous ARP requests for the router's IP address. This results in hosts sending their router traffic to the attacker, and the attacker forwarding that data to the router. This allows passwords, keys, and other information to be intercepted.

To protect against this type of attack, the router will send out its own gratuitous ARP request to override the attacker whenever a gratuitous ARP broadcast with the router's IP address as the source is received on the network.

This command disables gratuitous ARP protection.

Example

The following example disables gratuitous ARP protection for VLAN corp:

```
disable iparp gratuitous protect vlan corp
```

History

This command was first available in ExtremeXOS 11.2.

The *vlan_list* option was added in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable iparp refresh

```
disable iparp {vr vr_name} refresh
```

Description

Disables IP ARP to refresh its IP ARP entries before timing out.

Syntax Description

<i>vr_name</i>	Specifies a VR or VRF.
----------------	------------------------

Default

Enabled.

Usage Guidelines

The purpose of disabling ARP refresh is to reduce ARP traffic in a high node count Layer 2 switching only environment.

If you do not specify a VR or VRF, the command applies to the current VR context.

Example

The following example disables IP ARP refresh:

```
disable iparp refresh
```

History

This command was first available in ExtremeXOS 10.1.

The **vr** option was added in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable ipforwarding broadcast

```
disable ipforwarding broadcast [ {vlan} vlan_name | vlan vlan_list]
```

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.
<i>vlan_name</i>	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 example disables forwarding of IP broadcast traffic for a VLAN "accounting":

```
disable ipforwarding broadcast vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

The **ignore-broadcast** and **fast-direct-broadcast** keywords were added in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms that use the Edge, Advanced Edge, or Core license. For information on the licenses available for each platform, see the [Switch Engine 32.3 Feature License Requirements](#) document.

disable ipforwarding broadcast

```
disable ipforwarding broadcast [ {vlan} vlan_name | vlan vlan_list]
```

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.
<i>vlan_name</i>	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 example disables forwarding of IP broadcast traffic for a VLAN "accounting":

```
disable ipforwarding broadcast vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

The **ignore-broadcast** and **fast-direct-broadcast** keywords were added in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms that use the Edge, Advanced Edge, or Core license. For information on the licenses available for each platform, see the [Switch Engine 32.3 Feature License Requirements](#) document.

disable ipforwarding ipv6

```
disable ipforwarding ipv6 [ {vlan} vlan_name | vlan vlan_list] | tunnel  
    tunnel_name | vr vr_name}
```

Description

Disables routing for one or all interfaces. If no argument is provided, disables routing for all interfaces on the current VR or VRF.

Syntax Description

<i>vlan_name</i>	Specifies an IPv6 configured <u>VLAN</u> .
<i>vlan_list</i>	Specifies a VLAN list of IDs.
<i>tunnel_name</i>	Specifies an IPv6 tunnel.
<i>vr_name</i>	Specifies a VR or VRF.

Default

Disabled.

Usage Guidelines

When new IPv6 interfaces are added, IPv6 forwarding is disabled by default.

Example

The following example disables forwarding of IPv6 traffic for a VLAN "accounting":

```
disable ipforwarding ipv6 vlan accounting
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

disable ipmcforwarding ipv6

```
disable ipmcforwarding ipv6 {{vlan} name}
```

Description

Disables IPv6 multicast forwarding on a router interface.

Syntax Description

<i>name</i>	Specifies a <u>VLAN</u> name.
-------------	-------------------------------

Default

Disabled.

Usage Guidelines

If no options are specified, all configured IPv6 interfaces are affected. When new IPv6 interfaces are created, IPv6 multicast forwarding is disabled by default.

Disabling IPv6 multicast forwarding disables any Layer 3 IPv6 multicast routing for the streams coming to the interface.

Example

The following example disables IPv6 multicast forwarding on VLAN accounting:

```
disable ipmcforwarding ipv6 vlan accounting
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IPv6 multicast feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

disable ipmcforwarding

```
disable ipmcforwarding {vlan name}
```

Description

Disables IP multicast forwarding on a router interface.

Syntax Description

<i>name</i>	Specifies a <u>VLAN</u> name.
-------------	-------------------------------

Default

Disabled.

Usage Guidelines

If no options are specified, all configured IP interfaces are affected. When new IP interfaces are added, IP multicast forwarding is disabled by default.

IP forwarding must be enabled before enabling IP multicast forwarding.

Disabling IP multicast forwarding disables any Layer 3 multicast routing for the streams coming to the interface.

Example

The following example disables IP multicast forwarding on the VLAN accounting:

```
disable ipmcforwarding vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IPv4 multicast feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

NEW! disable ipmcforwarding tunnel

```
disable ipmcforwarding {tunnel tunnel_name}
```

Description

Disables IP multicast tunnel forwarding on a router interface.

Syntax Description

<i>tunnel_name</i>	Specifies a tunnel name.
--------------------	--------------------------

Default

Disabled.

Usage Guidelines

If no options are specified, all configured IP interfaces are affected. When new IP interfaces are added, IP multicast forwarding is disabled by default.

IP forwarding must be enabled before enabling IP multicast forwarding.

Disabling IP multicast forwarding disables any Layer 3 multicast routing for the streams coming to the interface.

Example

The following example disables IP multicast forwarding on the tunnel called accounting:

```
disable ipmcforwarding tunnel accounting
```

History

This command was first available in ExtremeXOS 32.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IPv4 multicast feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

disable ip option loose-source-route

```
disable ip option loose-source-route
```

Description

Disables processing of the loose source route IP option in the IPv4 packet header.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

Disables the switch from forwarding IP packets with the IP option for loose source routing turned on. Packets with the loose-source-route option enabled are dropped by the switch.

Example

The following example disables processing of the loose source route IP option:

```
# disable ip option loose-source-route
```

History

This command was first available in ExtremeXOS 10.1.

This command was removed in ExtremeXOS 30.1, and then re-introduced in ExtremeXOS 31.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable ip option strict-source-route

```
disable ip option strict-source-route
```

Description

Disables processing the strict source route IP option in the IPv4 packet header.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

Disables the switch from forwarding IP packets that have the strict source routing IP option turned on. The switch drops packets that have the strict source routing IP option enabled.

Example

The following example disables processing of the strict source route IP option:

```
# disable ip option strict-source-route
```

History

This command was first available in ExtremeXOS 10.1.

This command was removed in ExtremeXOS 30.1, and then re-introduced in ExtremeXOS 31.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable iproute bfd

```
disable iproute bfd {gateway} ip_addr {vr vrname}
```

Description

Disables BFD client services for IPv4 static routes.

Syntax Description

<i>ip_addr</i>	Specifies the IPv4 address of a neighbor for which BFD services are to be stopped.
<i>vrname</i>	Specifies the VR or VRF name for which BFD services are being disabled.

Default

Disabled.

Usage Guidelines

When the BFD client is disabled, BFD services for all static IP routes terminates. This command does not disable services for other BFD clients (such as the *MPLS* BFD client).

Example

The following example disables BFD client protection for communications with neighbor 10.10.10.1:

```
# disable iproute bfd 10.10.10.1
```

History

This command was first available in ExtremeXOS 12.5.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable iproute bfd strict

```
disable iproute {protection} bfd strict
```

Description

Turns off "strict" Bidirectional Forwarding Detection (BFD) session control, which brings up the static route during switch reboot if the static route nexthop BFD session is in the INIT state.

Syntax Description

protection	Enables or disables route protection.
bfd	BFD protect static routes to next hop gateway.
strict	Disables considering that protected static routes are not up if the BFD session is in INIT state. Default is disabled.

Default

By default, strict BFD session control is disabled.

Usage Guidelines

If the BFD session is down, but BFD protected static route is still in the routing table after reboot, the BFD session is never established, because during reboot, the BFD session is in the INIT state, and the static route is brought up without considering BFD session state. This can cause traffic loss because the link to the gateway actually is down. This command turns off strict BFD session control, which means that the static route is brought up during reboot even if the BFD session is in the INIT state. A reboot is required to make the command take effect.

Example

The following example disables BFD strict session control:

```
# disable iproute bfd strict
WARNING: Please reboot the switch for the strict BFD to take effect.
```

History

This command was first available in ExtremeXOS 30.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable iproute compression

disable iproute compression {**vr** *vrname*}

Description

Disables IPv4 route compression.

Syntax Description

<i>vrname</i>	VR or VRF name for which the IP route compression is being disabled. If the VR or VRF name is not specified, route compression is disabled for the VR context from which CLI command is issued.
---------------	---

Default

Enabled.

Usage Guidelines

Disables IPv4 route compression for a specified VR or VRF.

Example

The following example disables IP route compression:

```
disable iproute compression
```

History

This command was first available in ExtremeXOS 12.0.

Default changed to enabled in ExtremeXOS 15.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable iproute ipv6 compression

```
disable iproute ipv6 compression {vr vr_name}
```

Description

This command disables IPv6 route compression.

Syntax Description

<i>vr_name</i>	Specifies a VR or VRF. If not specified, the current CLI context is used.
----------------	---

Default

By default, IPv6 route compression is disabled for all address families and VRs.

Usage Guidelines

This command disables IPv6 route compression for the IPv6 address family and VR. This command decompresses previously compressed prefixes in the IPv6 prefix database.

Example

The following example disables IPv6 route compression for the IPv6 address family and the VR of the current CLI context:

```
disable iproute ipv6 compression
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

disable iproute ipv6 sharing

```
disable iproute ipv6 sharing {{{vr} vr_name} | {{{vr} all}}}
```

Description

This command disables IPv6 route sharing.

Syntax Description

<i>vr_name</i>	Specifies a VR or VRF. If not specified, the current CLI context is used
all	Specifies all VR or VRF.

Default

By default, IPv6 route sharing is disabled.

Usage Guidelines

This command disables IPv6 route sharing for the IPv6 address family and VR.

Example

The following example disables IPv6 route sharing for the IPv6 address family and the VR of the current CLI context:

```
disable iproute ipv6 sharing
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

The ability to enable and disable ECMP for IPv6 is supported for all platforms.

disable iproute mpls-next-hop

```
disable iproute mpls-next-hop
```

Description

Disables IP forwarding over *MPLS* LSPs for the default VR.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command disables IP forwarding over MPLS LSPs for the default VR. When disabled, any route with an MPLS LSP as its next hop becomes inactive and is not used to tunnel IP traffic across the MPLS network. By default, IP forwarding over MPLS LSPs is disabled.

Example

This command disables IP forwarding over MPLS LSPs.

```
disable iproute mpls-next-hop
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

disable iproute protection ping

```
disable iproute {ipv4 | ipv6} protection ping
```

Description

Globally disables ping protection for static routes added with ping protection for IPv4 and IPv6. Routes are up in the routing table, and ping health check monitoring is not performed.

Syntax Description

ipv4	Specifies IPv4 (default).
ipv6	Specifies IPv6.

protection	Disables route protection.
ping	Globally disables ping protection for static routes added with ping protection (default is enabled).

Default

Enabled is the default. If not specified, IPv4 is the default.

Example

The following example disables ping protection for static routes added with ping protection for IPv4:

```
# disable iproute ipv4 protection ping
```

History

This command was first available in ExtremeXOS 22.1.

Platform Availability

This command is available on all platforms with any license level as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

disable iproute sharing

```
disable iproute {ipv4} sharing {{vr} vrname} | {vr} {all}
```

Description

Disables IPv4 route sharing.

Syntax Description

vrname	VR or VRF name for which IP route sharing is being disabled.
---------------	--

Default

Disabled.

Usage Guidelines

If a VR is not specified, this command disables IP route sharing in the current VR context.

Example

The following example disables load sharing for multiple routes:

```
disable iproute sharing
```

History

This command was first available in ExtremeXOS 12.1.

The **vr** option was added in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable ip-security anomaly-protection icmp

```
disable ip-security anomaly-protection icmp {slot [ slot | all ]}
```

Description

Disables *ICMP* size and fragment checking.

Syntax Description

<i>slot</i>	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is disabled.

Usage Guidelines

This command disables ICMP size and fragment checking. This checking takes effect for both IPv4 and IPv6 TCP packets. When enabled, the switch drops ICMP packets if one of following condition is true:

- Fragmented ICMP packets for IPv4 packets.
- IPv4 ICMP pings packets with payload size greater than the maximum IPv4 ICMP-allowed size. (The maximum allowed size is configurable.)
- IPv6 ICMP ping packets with payload size > the maximum IPv6 ICMP-allowed size. (The maximum allowed size is configurable.)

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable ip-security anomaly-protection ip

```
disable ip-security anomaly-protection ip { slot [ slot | all ] }
```

Description

Disables source and destination IP address checking.

Syntax Description

<i>slot</i>	Specifies the slot.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is disabled.

Usage Guidelines

This command disables source and destination IP addresses checking. This checking takes effect for both IPv4 and IPv6 packets. When enabled, the switch drops IPv4/IPv6 packets if its source IP address are the same as the destination IP address. In most cases, the condition of source IP address being the same as the destination IP address indicates a Layer 3 protocol error. (These kind of errors are found in LAND attacks.)

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable ip-security anomaly-protection l4port

```
disable ip-security anomaly-protection l4port [tcp | udp | both] {slot  
[ slot | all ]}
```

Description

Disables TCP and UDP ports checking.

Syntax Description

tcp	Specifies that the TCP port be disabled for checking.
udp	Specifies that the UDP port be disabled for checking.

both	Specifies both the TCP and UDP ports be disabled for checking.
<i>slot</i>	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is disabled.

Usage Guidelines

This command disables TCP and UDP ports checking. This checking takes effect for both IPv4 and IPv6 TCP and UDP packets. When enabled, the switch drops TCP and UDP packets if its source port is the same as its destination port. In most cases, when the condition of source port is the same as that of the destination port, it indicates a Layer4 protocol error. (This type of error can be found in a BALT attack.)

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable ip-security anomaly-protection notify

```
disable ip-security anomaly-protection notify [log | snmp | cache] {slot
[ slot | all ]}
```

Description

Disables protocol anomaly notification.

Syntax Description

log	Specifies the switch to send the notification to a log file.
snmp	Specifies the switch to send an <i>SNMP</i> trap when an event occurs.
cache	Specifies the switch to send the notification to cache.
<i>slot</i>	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is disabled.

Usage Guidelines

This command disables anomaly notification. When enabled, any packet failed to pass enabled protocol checking is sent to XOS Host CPU and notifies the user. There are three different types of notifications:

- **log**: log anomaly events in the switch log system; you can view and manage this log with the `show log` and `configure log` commands.
- **snmp**: the anomaly events generate SNMP traps.
- **cache**: logs the most recent and unique anomaly events in memory; rebooting the switch will cause all the logged events to be lost (the number of cached events is configured by command).

When disabled, the switch drops all violating packets silently.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable ip-security anomaly-protection tcp flags

```
disable ip-security anomaly-protection tcp flags {slot [ slot | all ]}
```

Description

Disables TCP flag checking.

Syntax Description

<i>slot</i>	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is disabled.

Usage Guidelines

This command disables TCP flag checking. This checking takes effect for both IPv4 and IPv6 TCP packets. When enabled, the switch drops TCP packets if one of following condition is true:

- TCP SYN flag==1 and the source port<1024
- TCP control flag==0 and the sequence number==0
- TCP FIN, URG, and PSH bits are set, and the sequence number==0
- TCP SYN and FIN both are set.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable ip-security anomaly-protection tcp fragment

```
disable ip-security anomaly-protection tcp fragment {slot [ slot |  
all ] }
```

Description

Disables TCP fragment checking.

Syntax Description

<i>slot</i>	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is disabled.

Usage Guidelines

This command disables TCP fragment checking. This checking takes effect for IPv4/IPv6. When it is enabled, the switch drops TCP packets if one of following condition is true:

- For the first IPv4 TCP fragment (its IP offset field==0), if its TCP header is less than the minimum IPv4 TCP header allowed size.
- If its IP offset field==1 (for IPv4 only).

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable ip-security anomaly-protection

```
disable ip-security anomaly-protection {slot [ slot | all ] }
```

Description

Disables all anomaly checking options.

Syntax Description

<i>slot</i>	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is disabled.

Usage Guidelines

This commands disables all anomaly checking options, including IP address, UDP/TCP port, TCP flag and fragment, and *ICMP* anomaly checking.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable ip-security arp gratuitous-protection

```
disable ip-security arp gratuitous-protection [dynamic | {vlan}  
vlan_name | all ]
```

Description

Disables gratuitous ARP protection on one or all *VLANs* on the switch.

Syntax Description

all	Specifies all VLANs configured on the switch.
<i>vlan-name</i>	Specifies the VLAN.
dynamic	Configuration options for dynamically created VLANs.

Default

By default, gratuitous ARP protection is disabled.

Usage Guidelines

Beginning with ExtremeXOS 11.6, this command replaces the `disable iparp gratuitous protect vlan` command.

Hosts can launch man-in-the-middle attacks by sending out gratuitous ARP requests for the router's IP address. This results in hosts sending their router traffic to the attacker, and the attacker forwarding that data to the router. This allows passwords, keys, and other information to be intercepted.

To protect against this type of attack, the router will send out its own gratuitous ARP request to override the attacker whenever a gratuitous ARP broadcast with the router's IP address as the source is received on the network.

This command disables gratuitous ARP protection.

Example

The following command disables gratuitous ARP protection for VLAN corp:

```
disable ip-security arp gratuitous-protection vlan corp
```

History

This command was first available in ExtremeXOS 11.6.

Dynamic VLAN option added in ExtremeXOS 30.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable ip-security arp learning learn-from-arp

```
disable ip-security arp learning learn-from-arp [dynamic | {vlan}  
vlan_name] ports [all | ports]
```

Description

Disables ARP learning on the specified VLAN and member ports.

Syntax Description

<i>vlan_name</i>	Specifies the name of the VLAN to which this rule applies.
dynamic	Configuration options for dynamically created VLANs.
all	Specifies all ingress ports.
<i>ports</i>	Specifies one or more ingress ports.

Default

By default, ARP learning is enabled.

Usage Guidelines

You can disable ARP learning so that the only entries in the ARP table are either manually added or those created by *DHCP* secured ARP; the switch does not add entries by tracking ARP requests and replies. By disabling ARP learning and adding a permanent entry or configuring DHCP secured ARP, you can centrally manage and allocate client IP addresses and prevent duplicate IP addresses from interrupting network operation.

To manually add a permanent entry to the ARP table, use the following command:

```
configure iparp add ip_addr {vrvr_name} mac
```

To configure DHCP secure ARP as a method to add entries to the ARP table, use the following command:

```
enable ip-security arp learning learn-from-dhcp vlan vlan_name  
ports [all | ports] {poll-interval interval_in_seconds} {retries  
number_of_retries}
```

Displaying ARP Information

To display how the switch builds an ARP table and learns MAC addresses for devices on a specific VLAN and associated member ports, use the following command:

```
show ip-security arp learning {vlan} vlan_name
```

To view the ARP table, including permanent and DHCP secured ARP entries, use the following command:

```
show iparp {ip_address | mac | vlanvlan_name | permanent} {vrvr_name}
```



Note

DHCP secured ARP entries are stored as static entries in the ARP table.

Example

The following command disables ARP learning on port 1:1 of the VLAN learn:

```
disable ip-security arp learning learn-from-arp vlan learn ports 1:1
```

History

This command was first available in ExtremeXOS 11.6.

Dynamic VLAN option was added in ExtremeXOS 30.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable ip-security arp learning learn-from-dhcp

```
disable ip-security arp learning learn-from-dhcp [dynamic vlan | {vlan}  
          vlan_name ports [all | ports]
```

Description

Disables *DHCP* secured ARP learning for the specified *VLAN* and member ports.

Syntax Description

dynamic	Configuration options for dynamically created VLANs.
<i>vlan_name</i>	Specifies the name of the VLAN to which this rule applies.
all	Specifies all ingress ports.
<i>ports</i>	Specifies one or more ingress ports.

Default

By default, DHCP secured ARP learning is disabled.

Usage Guidelines

Use this command to disable DHCP secured ARP learning.

Displaying ARP Information

To display how the switch builds an ARP table and learns MAC addresses for devices on a specific VLAN and associated member ports, use the following command:

```
show ip-security arp learning {vlan} vlan_name
```

To view the ARP table, including permanent and DHCP secured ARP entries, use the following command:

```
show iparp {ip_address | mac | vlanvlan_name | permanent} {vrr_name}
```

Example

The following command disables DHCP secured ARP learning on port 1:1 of the VLAN learn:

```
disable ip-security arp learning learn-from-dhcp vlan learn ports 1:1
```

History

This command was first available in ExtremeXOS 11.6.

Dynamic VLAN support was added in ExtremeXOS 30.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable ip-security arp validation

```
disable ip-security arp validation [dynamic | {vlan} vlan_name] [all |  
    ports]
```

Description

Disables ARP validation for the specified *VLAN* and member ports.

Syntax Description

<i>vlan_name</i>	Specifies the name of the VLAN to which this rule applies.
dynamic	Configuration options for dynamically created VLANs.
all	Specifies all ports.
<i>ports</i>	Specifies one or more ports.

Default

By default, ARP validation is disabled.

Usage Guidelines

Use this command to disable ARP validation.

Displaying ARP Validation Information

To display information about ARP validation, use the following command:

```
show ip-security arp validation {vlan} vlan_name
```

Example

The following command disables ARP validation on port 1:1 of the VLAN valid:

```
disable ip-security arp validation vlan valid ports 1:1
```

History

This command was first available in ExtremeXOS 11.6.

Dynamic VLAN option added in ExtremeXOS 30.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable ip-security dhcp-bindings restoration

```
disable ip-security dhcp-bindings restoration
```

Description

Disables the download and upload of *DHCP* bindings.

Syntax

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

The command allows you to disable the download and upload of the DHCP bindings, essentially disabling the DHCP binding functionality. The default is disabled.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable ip-security dhcp-snooping

```
disable ip-security dhcp-snooping [dynamic | {vlan} vlan_name] ports
[all | ports]
```

Description

Disables *DHCP* snooping on the switch.

Syntax Description

<i>vlan_name</i>	Specifies the name of the DHCP-snooping <i>VLAN</i> .
dynamic	Configuration options for dynamically created VLANs.

all	Specifies all ports to stop receiving DHCP packets.
<i>ports</i>	Specifies one or more ports to stop receiving DHCP packets.

Default

By default, DHCP snooping is disabled.

Usage Guidelines

Use this command to disable DHCP snooping on the switch.

Example

The following command disables DHCP snooping on the switch:

```
disable ip-security dhcp-snooping vlan snoop ports 1:1
```

History

This command was first available in ExtremeXOS 11.6.

Dynamic VLAN option added in ExtremeXOS 30.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable ip-security source-ip-lockdown ports

```
disable ip-security source-ip-lockdown ports [all | ports]
```

Description

Disables the source IP lockdown feature on one or more ports.

Syntax Description

all	Specifies all ports for which source IP lockdown should be disabled.
<i>ports</i>	Specifies one or more ports for which source IP lockdown should be disabled.

Default

By default, source IP lockdown is disabled on the switch.

Usage Guidelines

To display the source IP lockdown configuration on the switch, use the following command:

```
show ip-security source-ip-lockdown
```

Example

The following command disables source IP lockdown on ports 1:1 and 1:4:

```
disable ip-security source-ip-lockdown ports 1:1, 1:4
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable iqagent

```
disable iqagent
```

Description

Disables the ExtremeCloud™ IQ Agent.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Disabling IQ Agent prevents all access to ExtremeCloud IQ. Any current activity with ExtremeCloud IQ, including remote SSH sessions, are disconnected immediately. Re-enabling IQ Agent can only occur by using the `enable iqagent` command using either console or Telnet or SSH access. Disabling IQ Agent deactivates automatic DHCP access on VLAN Mgmt, which is required for Zero-Touch Provisioning (ZTP).

To view the state if the IQ Agent, use the command `show iqagent discovery` without the **discovery** option.

Example

The following example disables the IQ Agent:

```
# disables iqagent
```

History

This command was first available in ExtremeXOS 31.1.

Platform Availability

This command is available on the ExtremeSwitching 5520 platform.

disable irdp

```
disable irdp {vlan name}
```

Description

Disables the generation of *ICMP* router advertisement messages on one or all *VLANs*.

Syntax Description

<i>name</i>	Specifies a VLAN name.
-------------	------------------------

Default

Disabled.

Usage Guidelines

If no optional argument is specified, all the IP interfaces are affected.

Example

The following example disables IRDP on VLAN "accounting":

```
disable irdp vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms that use the Edge, Advanced Edge, or Core license. For information on the licenses available for each platform, see the [Switch Engine 32.3 Feature License Requirements](#) document.

disable isis

```
disable isis {area area_name}
```

Description

This command disables the specified IS-IS router process on the current virtual router.

Syntax Description

<i>area_name</i>	Specifies the name of the IS-IS router process to be disabled.
------------------	--

Default

Disabled.

Usage Guidelines

IS-IS PDUs are no longer sent or processed on this IS-IS router process. The LSP and neighbor databases are purged. IS-IS routes are purged from the routing table. This command should only be used during planned network outages. This command has no effect on router processes that are already disabled.

Example

The following command disables the IS-IS process named areax:

```
disable isis area areax
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Premier license as described in the [**Switch Engine 32.3 Feature License Requirements**](#) document.

disable isis area adjacency-check

```
disable isis area area_name adjacency-check {ipv4 | ipv6}
```

Description

This command disables the checking of the following TLVs when forming adjacencies: Protocols Supported and IP Interface Address.

Syntax Description

<i>area_name</i>	Specifies the name of the IS-IS router process that should no longer perform the adjacency check.
ipv4	Specifies that the adjacency check should no longer be performed on IPv4 interfaces.
ipv6	Specifies that the adjacency check should no longer be performed on IPv6 interfaces.

Default

IPv4: Enabled.

IPv6: Enabled.

Usage Guidelines

When adjacency checking is disabled, adjacencies may be formed on interfaces that do not reside on the same subnet or do not support IPv4 (if disabled for IPv4) or IPv6 (if disabled for IPv6). If neither `ipv4` nor `ipv6` is specified, this command applies to IPv4.

Example

The following command directs the IS-IS process named `areax` to disable adjacency checks on IPv6 interfaces:

```
disable isis area areax adjacency-check ipv6
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Premier license as described in the [***Switch Engine 32.3 Feature License Requirements***](#) document.

disable isis area dynamic-hostname

```
disable isis area area_name dynamic-hostname
```

Description

This command disables the dynamic hostname feature.

Syntax Description

<i>area_name</i>	Specifies the name of the IS-IS process for which the dynamic-hostname feature is to be disabled.
------------------	---

Default

Disabled.

Usage Guidelines

The specified router process no longer includes code 137 TLVs in its LSPs and names are no longer displayed in show commands.

Example

The following command disables the display of area names or *SNMP* names instead of system IDs:

```
disable isis area areax dynamic-hostname
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Premier license as described in the *Switch Engine 32.3 Feature License Requirements* document.

disable isis area export ipv6

```
disable isis area area_name export ipv6 route-type
```

Description

This command disables IPv6 route redistribution of the specified type into IS-IS.

Syntax Description

<i>area_name</i>	Specifies the IS-IS router process for which route redistribution is disabled.
<i>route-type</i>	Selects the type of export route to disable. The valid route types are: direct, ospfv3, ospfv3-extern1, ospfv3-extern2, ospfv3-inter, ospfv3-intra, ripng, bgp, and static.

Default

All types are disabled.

Usage Guidelines

None.

Example

The following command disables *RIPng* route distribution into areax:

```
disable isis area areax export ipv6 ripng
```

History

This command was first available in ExtremeXOS 12.1.
Support for *BGP* was added in ExtremeXOS 12.6.0-BGP.

Platform Availability

This command is available on platforms with a Premier license as described in the *Switch Engine 32.3 Feature License Requirements* document.

disable isis area export

```
disable isis area area_name export {ipv4} route-type
```

Description

This command disables IPv4 route redistribution of the specified type into IS-IS.

Syntax Description

<i>area_name</i>	Specifies the IS-IS router process for which route redistribution is disabled.
ipv4	Specifies that the configuration change is for IPv4 IS-IS routing.
<i>route-type</i>	Selects the type of export route to disable. The valid route types are: bgp, direct, e-bgp, i-bgp, ospf, ospf-extern1, ospf-extern2, ospf-inter, ospf-intra, rip, and static.

Default

All types are disabled.

Usage Guidelines

None.

Example

The following command disables *RIP* route distribution into areax:

```
disable isis area areax export rip
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Premier license as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

disable isis area originate-default

```
disable isis area area_name originate-default {ipv4 | ipv6}
```

Description

This command disables the generation of one or all default routes in the LSPs for the specified router process.

Syntax Description

<i>area_name</i>	Specifies the name of the IS-IS router process that should no longer generate the default route.
ipv4	Specifies that the router process should no longer generate the default IPv4 route.
ipv6	Specifies that the router process should no longer generate the default IPv6 route.

Default

IPv4: Disabled.

IPv6: Disabled.

Usage Guidelines

This applies to level 2 routing only. By default this command disables IPv4 default route origination. The optional ipv6 keyword disables IPv6 default route origination. This command has no effect on router processes that are already disabled for default route origination on level 1-only router processes.

Example

The following command directs the IS-IS process named areax to stop generating the default IPv4 route in it's LSPs:

```
disable isis area areax originate-default
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Premier license as described in the [***Switch Engine 32.3 Feature License Requirements***](#) document.

disable isis area overload-bit

```
disable isis area area_name overload-bit
```

Description

This command disables the overload-bit feature.

Syntax Description

<i>area_name</i>	Specifies the area name of the IS-IS process for which this feature is to be disabled.
------------------	--

Default

Disabled.

Usage Guidelines

Disabling the overload bit feature causes an SPF recalculation throughout the network. In addition, external and interlevel router redistribution is no longer suppressed if those options were included when the overload bit was enabled. If the overload bit is currently set as a result of the overload-bit on-startup command, this command overrides the configuration and disables this feature.

Example

The following command disables the overload bit feature for areax:

```
disable isis area areax overload-bit
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Premier license as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

disable isis hello-padding

```
disable isis [vlan all | {vlan} vlan_name] hello-padding
```

Description

This command disables the padding of Hello PDUs for one or all IS-IS [VLANs](#).

Syntax Description

vlan all	Disables hello padding on all IS-IS VLANs.
vlan_name	Specifies a single VLAN on which to disable hello padding.

Default

Enabled.

Usage Guidelines

Implicit adjacency MTU verification is not performed when hello padding is disabled.

Example

The following command disables hello padding on all IS-IS VLANs:

```
disable isis vlan all hello-padding
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Premier license as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

disable isis restart-helper

```
disable isis restart-helper
```

Description

This command disables the IS-IS restart helper.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

When this feature is disabled, the router does not act as a restart helper and may time out a restarting router's adjacency per normal operation.

Example

The following command disables the IS-IS restart helper:

```
disable isis restart-helper
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Premier license as described in the [***Switch Engine 32.3 Feature License Requirements***](#) document.

disable jumbo-frame ports

```
disable jumbo-frame ports [all | port_list]
```

Description

Disables jumbo frame support on a port.

Syntax Description

all	Specifies all ports.
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

Disabled.

Usage Guidelines

Use this command to disable jumbo frames on individual ports.

Example

The following command disables jumbo frame support on a switch:

```
disable jumbo-frame ports all
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable l2vpn

```
disable l2vpn [vppls [vppls_name | all] | vpws [vpws_name | all]]
```

Description

Disables the specified VPLS or VPWS.

Syntax Description

<i>vppls_name</i>	Identifies the VPLS within the switch (character string).
<i>vpws_name</i>	Identifies the VPWS within the switch (character string).
all	Specifies all VPLS or VPWS instances.

Default

All newly created VPLS instances are enabled.

Usage Guidelines

When a VPLS or VPWS instance is disabled, all sessions to its configured peers are terminated. Any locally attached service [VLAN/VMAN](#) is immediately isolated from other devices residing in the VPN. If this is an H-VPLS core node, then all spoke nodes connected to this peer are isolated unless redundant core access is configured.

The **l2vpn** keyword is introduced in ExtremeXOS Release 12.4 and is required when disabling a VPWS. For backward compatibility, the **l2vpn** keyword is optional when disabling a VPLS. However, this keyword will be required in a future release, so we recommend that you use this keyword for new configurations and scripts.

Example

The following command disables the VPLS named myvpls:

```
disable vpls myvpls
```

The following command disables the VPWS named myvpws:

```
disable l2vpn vpws myvpws
```

History

This command was first available in ExtremeXOS 11.6.

The **l2vpn** and **vpws** keywords were first available in ExtremeXOS 12.4.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

disable l2vpn health-check vccv

```
disable l2vpn [vpls [vpls_name | all] | vpws [vpws_name | all]] health-check vccv
```

Description

Disables the VCCV health check feature on the specified VPLS or VPWS instances.

Syntax Description

<i>vpls_name</i>	Identifies the VPLS for which health check is to be disabled.
<i>vpws_name</i>	Identifies the VPWS for which health check is to be disabled.
all	Specifies that health check is to be disabled on all VPLS instances on the local node.

Default

Health check is disabled.

Usage Guidelines

The **l2vpn** keyword was introduced in ExtremeXOS Release 12.4 and is required when disabling health check for a VPWS instance. For backward compatibility, the **l2vpn** keyword is optional when disabling health check for VPLS instance. However, this keyword will be required in a future release, so we recommend that you use this keyword for new configurations and scripts.

Example

The following command disables the health check feature on the VPLS instance myvpls:

```
disable l2vpn vpls myvpls health-check vccv
```

History

This command was first available in ExtremeXOS 12.1.

The **l2vpn** and **vpws** keywords were first available in ExtremeXOS 12.4.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

disable l2vpn service

```
disable l2vpn [vpls [vpls_name | all] | vpws [vpws_name | all]] service
```

Description

Disables the configured services for the specified VPLS or VPWS.

Syntax Description

<i>vpls_name</i>	Identifies the VPLS within the switch (character string).
<i>vpws_name</i>	Identifies the VPWS within the switch (character string).
all	Specifies all VPLS or VPWS instances.

Default

Enabled.

Usage Guidelines

When services are disabled, the VPLS or VPWS is removed from all peer sessions. The keyword all disables services for all VPLS instances.

The **l2vpn** keyword was introduced in ExtremeXOS Release 12.4 and is required when disabling a service for a VPWS peer. For backward compatibility, the **l2vpn** keyword is optional when disabling a service for a VPLS peer. However, this keyword will be required in a future release, so we recommend that you use this keyword for new configurations and scripts.

Example

The following command disables the configured services for VPLS myvpls:

```
disable l2vpn vpls myvpls service
```

History

This command was first available in ExtremeXOS 11.6.

The **l2vpn** and **vpws** keywords were first available in ExtremeXOS 12.4.

Platform Availability

This command is available only on the platforms that support *MPLS* as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

disable l2vpn sharing

```
disable l2vpn sharing
```

Description

Disables LSP sharing for Layer 2 VPN pseudo-wires.

Syntax Description

This command has no keywords or arguments.

Default

Disabled.

Usage Guidelines

This command disables LSP sharing for L2VPN PWs. When LSP sharing is disabled, only 1 named LSP is used for a PW. When LSP sharing is enabled, up to 16 named LSPs are used for a PW.

If LSP Sharing is disabled, and more than 1 Transport LSP is programmed into HW, all but 1 Transport LSP is removed from HW, and the configuration is preserved. If LSP Sharing is enabled, and more than 1 Transport LSP was previously configured, the remaining LSPs is programmed into HW as they become available for use.

Example

The following command disables LSP sharing for L2VPN PWs:

```
disable l2vpn sharing
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available only on the platforms that support as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

disable l2vpn vpls peer fdb send-mac-withdrawal

```
disable l2vpn vpls peer [ipaddress | all] fdb send-mac-withdrawal
```

Description

Disables the MAC address withdrawal capability.

Syntax Description

l2vpn	Designates L2 VPN configuration.
vpls	Designates VPLS of MPLS configuration.
peer	Designates VPLS peer.
<i>ipaddress</i>	Selects the VPLS peer of the provided IP address.
all	Selects all VPLS peers.
fdb	Designates FDB.
send-mac-withdrawal	Disables sending the MAC address withdrawal message.

Default

Enabled.

Usage Guidelines

When disabled, the switch does not send MAC address withdrawal messages. If a MAC address withdrawal message is received from another VPLS peer, the local peer processes the message and withdraws the specified MAC addresses from its *FDB*, regardless of the MAC address withdrawal configuration.

Example

The following command disables MAC address withdrawal message for all VPLS peers:

```
# disable l2vpn vpls peer all fdb send-mac-withdrawal
```

History

This command was first available in ExtremeXOS 12.1.

The **l2vpn** keyword was added in ExtremeXOS 12.4.

Platform Availability

This command is available only on the platforms that support *MPLS* as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

disable learning iparp sender-mac

```
disable learning iparp {vr vr_name} sender-mac
```

Description

Disables MAC address learning from the payload of IP ARP packets.

Syntax Description

<i>vr_name</i>	Specifies a virtual router.
----------------	-----------------------------

Default

Disabled.

Usage Guidelines

To view the configuration for this feature, use the following command: `show iparp`

Example

The following example disables MAC address learning from the payload of IP ARP packets:

```
disable learning iparp sender-mac
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable learning port

```
disable learning {drop-packets | forward-packets} port [port_list | all]
```

Description

Disables MAC address learning on one or more ports for security purposes.

Syntax Description

drop-packets	Specifies that packets with unknown source MAC addresses be dropped. When disable learning is configured, this is the default behavior.
forward-packets	Specifies that packets with unknown source MAC addresses be forwarded.
port	Specifies the port.
<i>port_list</i>	Specifies one or more ports or slots and ports.
all	Specifies all ports and slots.

Default

Enabled.

Usage Guidelines

Use this command in a secure environment where access is granted via permanent forwarding databases (FDBs) per port.

Example

The following command disables MAC address learning on port 4 on a switch:

```
disable learning ports 4
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable learning vxlan ipaddress

```
disable learning {forward-packets | drop-packets} vxlan {vr vr_name}  
ipaddress remote_ipaddress
```

Description

This command disables learning a remote endpoint.

Syntax Description

forward-packets	Forward packets with unknown source MAC addresses.
drop-packets	Drop packets with unknown source MAC addresses.
vr	VR/VRF instance the IPv4 address is configured on.
<i>vr_name</i>	An existing VR/VRF name.

Default

N/A.

Usage Guidelines

N/A.

Example

To disable learning on a remote endpoint:

```
disable learning vxlan ipaddress 1.2.3.4
```

History

This command was first available in ExtremeXOS 21.1.

Platform Availability

This command is supported on the ExtremeSwitching 5320, 5420, 5520, and 5720 series switches, and stacks with 5320, 5420, 5520, and 5720 slots only.

disable led locator

```
disable led locator { slot [slot | all ] }
```

Description

Disables the front panel LEDs from flashing on a switch.

Syntax Description

slot <i>slot</i>	Slot number.
all	All slots.

Default

N/A.

Usage Guidelines

None.

Example

The following example disables the front panel LEDs on all slots:

```
disable led locator all
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable lldp ports

```
disable lldp ports [all | port_list] {receive-only | transmit-only}
```

Description

Disables LLDP transmit mode, receive mode, or transmit and receive mode on the specified port or ports.

Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.
receive-only	Specifies that only the receive mode for LLDP is disabled.
transmit-only	Specifies that only the transmit mode for LLDP is disabled.

Default

Enabled.

Usage Guidelines

If you do not specify an option, both LLDP modes (transmit and receive) are disabled.

Example

The following example disables the LLDP receive mode on ports 1:2 to 1:6.

```
disable lldp ports 1:2-1:6 receive-only
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable log debug-mode

```
disable log debug-mode
```

Description

Disables debug mode. The switch stops generating debug events.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command disables debug mode. Debug mode must be enabled prior to configuring advanced debugging capabilities. These include allowing debug messages, which can severely degrade performance. For typical network device monitoring, debug mode should remain disabled, the default setting. Debug mode should only be enabled when advised by technical support, or when advanced diagnosis is required. The debug mode setting is saved to FLASH.

The following configuration options require that debug mode be enabled:

- Including a severity of debug-summary, debug-verbose, or debug-data when configuring filters.

- Target format options process-name, process-id, source-function, and source-line.

Example

The following command disables debug mode:

```
disable log debug-mode
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable log display

```
disable log display
```

Description

Disables the sending of messages to the console display.

In a stack, this command is applicable only to Master and Backup nodes and not applicable to the standby nodes.

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.

This command setting is saved to FLASH and determines the initial setting of the console display at boot up.

You can also use this following command to control logging to different targets:

This command is equivalent to `disable log target console-display` command.

Example

The following command disables the log display:

```
disable log display
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable log target

```
disable log target [console | memory-buffer | nvram | primary-node |  
  backup-node | session | syslog [all | ipaddress udp-port {udp_port}  
  | ipPort | ipaddress tls_port {tls_port}] {vr vr_name} {local0 ...  
  local17}]]
```

Description

Stops sending log messages to the specified target.

In a stack, this command is applicable only to Master and Backup nodes and not applicable to the standby nodes.

Syntax Description

console	Specifies the console display.
memory-buffer	Specifies the switch memory buffer.
nvr am	Specifies the switch NVRAM.
primary-node	Specifies the primary node in a stack.
backup-node	Specifies the backup node in a stack.
session	Specifies the current session (including console display).
syslog	Specifies a syslog target.
all	Specifies all of the remote syslog servers.
<i>ipaddress</i>	Specifies the syslog host name or IP address.
<i>ipPort</i>	Specifies the UDP port number for the syslog target.
tls_port	Specifies remote Syslog server Transport Layer Security (TLS) for connection type.
<i>tls_port</i>	TLS port number.

<i>vr_name</i>	Specifies the virtual router that can reach the server IP address. Note: User-created VRs are supported only on the platforms listed for this feature in the Switch Engine 32.3 Feature License Requirements document..
local0 ... local7	Specifies the local syslog facility.

Default

Enabled, for memory buffer, NVRAM, primary node, and backup node; all other targets are disabled by default.

Usage Guidelines

This command stops sending messages to the specified target. By default, the memory buffer, NVRAM, primary node, and backup node targets are enabled. Other targets must be enabled before messages are sent to those targets.

Configuration changes to the session target are in effect only for the duration of the console display or Telnet session, and are not saved in FLASH. Changes to the other targets are saved to FLASH.

You can also use the following command to disable displaying the log on the console: `disable log display`

The `disable log display` command is equivalent to disable log target console-display command.

Example

The following example disables log messages to the current session:

```
disable log target session
```

History

- This command was first available in ExtremeXOS 10.1.
- The **ipPort** parameter was first available in ExtremeXOS 11.0.
- The udp port parameter was added in ExtremeXOS 21.1.
- Transport Layer Security (TLS) option added in ExtremeXOS 22.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable log target upm

```
disable log target upm {upm_profile_name}
```

Description

Disables the specified UPM log target.

Syntax Description

<i>upm_profile_name</i>	Specifies the name of the UPM log target to be disabled.
-------------------------	--

Default

N/A.

Usage Guidelines

This command disables the log target and retains any configurations applied to that target. To delete a target and any configuration applied to the target, use the following command:

```
delete log target upm {upm_profile_name}
```

Example

The following example disables the UPM log target "testprofile1":

```
disable log target upm testprofile1
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the Universal Port feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

disable log target xml-notification

```
disable log target xml-notification xml_target_name
```

Description

Disables a Web server target.

Syntax Description

<i>xml_target_name</i>	Specifies the name of the xml-notification target.
------------------------	--

Default

N/A.

Usage Guidelines

Use this command to disable a web server EMS target.

Example

The following command disables the Web server target target2:

```
disable log target xml-notification target2
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable loopback-mode vlan

disable loopback-mode [{**vlan**} *vlan_name* | **vlan** *vlan_list*]

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

<i>vlan_name</i>	Specifies a VLAN name.
<i>vlan_list</i>	Specifies a VLAN list of IDs.

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 example 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 ExtremeXOS 10.1.

The *vlan_list* option was added in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable mac-lockdown-timeout ports

```
disable mac-lockdown-timeout ports [all | port_list]
```

Description

Disables the MAC address lock down timeout feature for the specified port or group of ports or for all ports on the switch.

Syntax Description

all	Specifies all ports.
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

By default, the MAC address lock down feature is disabled.

Usage Guidelines

If you disable the MAC lock down timer on a port, existing MAC address entries for the port will time out based on the *FDB* aging period.

Example

The following command disables the MAC address lock down timer set for ports 2:3 and 2:4:

```
disable mac-lockdown-timeout ports 2:3, 2:4
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable mac-locking ports

```
disable mac-locking ports [port_list | all]
```

Description

Disables MAC locking on the specified port.

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
all	Specifies all ports.

Default

MAC locking is disabled by default.

Usage Guidelines

None.

Example

The following example disables MAC locking on port 14:

```
disable mac-locking ports 14
```

History

This command was first available in ExtremeXOS 15.7.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable mac-locking

```
disable mac-locking
```


Description

Disables MAC locking globally on the switch.

Syntax Description

This command has no arguments or variables.

Default

MAC locking is disabled by default.

Usage Guidelines

If you disable MAC locking globally, you cannot enable MAC locking on a specific port.

Example

The following example disables MAC locking on the switch.

```
disable mac-locking
```

History

This command was first available in ExtremeXOS 15.7.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable mirror

```
disable mirror [mirror_name | all]
```

Description

Disables a mirror instance.

Syntax Description

<i>mirror_name</i>	Specifies the mirror name.
all	Specifies all mirror instance are deleted.

Default

Disabled.

Usage Guidelines

Use this command to disable mirrors. Disabling an instance only changes the state, its configuration remains as defined (a change from current operation, which loses some configuration parameters).

Example

The following example disable a mirror instance named "mirror1" :

```
disable mirror mirror1
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable mirror control_index

```
disable mirror control_index {mirror mirror_name}
```

Description

Disables a Mirror MIB instance or the assigned instance to an existing mirror.

Syntax Description

<i>control_index</i>	Selects the Mirror MIB instance to disable. Range is 1 through 4.
mirror	Designates specifying a mirror name associated within the specified control index.
<i>mirror_name</i>	Specifies the mirror name associated within the specified control index.

Default

Disabled.

Usage Guidelines

Specifying a mirror name only disables that mirror within the Mirror MIB group (control index).

Example

The following example disables Mirror MIB specified by control index "1":

```
# disable mirror 1
```

The following example disables the mirror named "m1" within the Mirror MIB specified by control index "1":

```
# disable 1 mirror m1
```

History

This command was first available in ExtremeXOS 30.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable mlag port

disable mlag port *port*

Description

Removes a local port or LAG from an .

Syntax Description

<i>port</i>	Specifies a local member port of the MLAG group.
-------------	--

Default

N/A.

Usage Guidelines

Use this command to remove a local port or LAG from an MLAG.

Example

The following command unbinds the local member port 2:

```
# disable mlag port 2
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, 5720 series switches.

disable mlag port reload-delay

```
disable mlag port reload-delay
```

Description

This command disables reload-delay on Multi-switch Link Aggregation Group (MLAG) ports.

Syntax Description

This command has no arguments or variables.

Default

MLAG reload-delay is disabled by default.

Usage Guidelines

There are cases where MLAG ports comes up quicker than ISC ports after a switch reboot causing traffic loss during this time gap. This command disables this timer feature.

Example

The following example disables the MLAG reload-delay timer:

```
# disable mlag port reload-delay
```

History

This command was first available in ExtremeXOS 22.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, 5720 series switches.

disable mld

```
disable mld {vlan name}
```

Description

Disables MLD on a router interface. If no VLAN is specified, MLD is disabled on all router interfaces.

Syntax Description

<i>name</i>	Specifies a VLAN name.
-------------	------------------------

Default

Disabled.

Usage Guidelines

MLD is a protocol used by an IPv6 host to register its IPv6 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, hosts respond to the query, and group registration is maintained.

MLD is disabled by default on the switch. However, the switch can be configured to enable the generation and processing of MLD packets. MLD should be enabled when the switch is configured to perform IPv6 unicast or IPv6 multicast routing.

This command disables all MLD versions. When MLD is disabled, the MLDv2 compatibility mode setting is lost. If compatibility mode is not specified in the command when MLD is enabled again, MLDv1 compatibility mode is set.

Example

The following example disables MLD on VLAN accounting:

```
disable mld vlan accounting
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 multicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

disable mld snooping

```
disable mld snooping {with-proxy | vlan name}
```

Description

Disables MLD snooping.

Syntax Description

with-proxy	Disables the MLD snooping proxy.
<i>name</i>	Specifies a VLAN .

Default

The **with-proxy** option is enabled by default.

Usage Guidelines

If a VLAN is specified, MLD snooping is disabled only on that VLAN, otherwise MLD snooping is disabled on all VLANs.

The with-proxy option can be used for troubleshooting purpose. It should be enabled for normal network operation.

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 MLD done messages so that they are forwarded only when the last member leaves the group.

Example

The following example disables MLD snooping on the VLAN accounting:

```
disable mld snooping accounting
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 multicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

disable mld-ssm map

```
disable mld-ssm map {{vr} vr_name}
```

Description

Disables MLD SSM mapping on a VR.

Syntax Description

vr vr_name	Specifies a virtual router name.
-------------------	----------------------------------

Default

Disabled.

Usage Guidelines

Use this command to disable MLD SSM mapping on a VR.

Example

The following example disables SSM mapping on VR1:

```
disable mld-ssm map vr vrl
```

History

This command was first available in ExtremeXOS 15.5.

Platform Availability

This command is available on the platforms listed for the IPv6 multicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

disable mpls

```
disable mpls
```

Description

Disables MPLS on the switch.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

When MPLS is disabled, no label traffic is received or transmitted, and all MPLS-related protocol peer sessions are terminated.

Example

The following command globally disables MPLS on the switch:

```
disable mpls
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

disable mpls bfd

```
disable mpls bfd [vlan all | {vlan} vlan_name] {delete-sessions}
```

Description

Disables the Bidirectional Forwarding Detection (BFD) client for *MPLS* on the specified *VLAN* or on all VLANs.

Syntax Description

<i>vlan_name</i>	Specifies the VLAN on which to disable the MPLS BFD client.
delete-sessions	Specifies to delete all MPLS BFD sessions.

Default

Keep existing MPLS BFD sessions.

Usage Guidelines

This command instructs MPLS to cease the establishment of new BFD sessions with neighbors as LSPs are established with those neighbors. The default behavior retains the existing BFD sessions and ignores status updates from those existing sessions. The **delete-sessions** option instructs MPLS to request the deletion of existing sessions. Whether the sessions are deleted or not, the link state presented to the upper MPLS layers reverts to the normal link operational status.



Note

Deleting existing sessions can result in a neighbor DOWN indication from BFD to MPLS on the other end of the session (the peer switch) and a subsequent interface DOWN indication presented to the upper layers of MPLS on that peer switch. These actions can cause MPLS to reroute or fail the affected LSPs.

To disable the MPLS BFD client and delete all BFD sessions without disrupting the LSPs between two switches, do the following:

- Log into switch A as an admin user and issue the command: `disable mpls bfd vlanx`.
- Log into switch B as an admin user and issue the command: `disable mpls bfd vlanx delete-sessions`

Example

The following command disables the MPLS BFD client on VLAN vlan1:

```
disable mpls bfd vlan1
```


History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

disable mpls exp examination

```
disable mpls exp examination
```

Description

Disables assigning an [MPLS](#) packet to a [QoS](#) profile based on the MPLS packet's EXP value.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command disables assigning an MPLS packet to a QoS profile based on the MPLS packet's EXP value.

When disabled, all received MPLS packets are assigned to QoS profile **qp1**.

Example

The following command disables the assignment of an MPLS packet to a QoS profile based on the MPLS packet's EXP value:

```
disable mpls exp examination
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

disable mpls exp replacement

```
disable mpls exp replacement
```

Description

Disables setting an *MPLS* packet's EXP value based on the packet's *QoS* profile.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command disables setting an MPLS packet's EXP value based on the packet's QoS profile. The QoS profiles to EXP value mappings are configured using the `configure mpls exp replacement` command.

When disabled, all MPLS packets are transmitted with an EXP value of zero.

Example

The following command disables the setting of an MPLS packet's EXP value based on the packet's QoS profile:

```
disable mpls exp replacement
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

disable mpls ldp bgp-routes

```
disable mpls ldp bgp-routes
```

Description

Disables LDP's use of IP prefixes learned from *BGP* when establishing LDP LSPs.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

This command disables LDP's establishment of LSPs to routes learned via BGP, thus reducing the internal resources used by LDP. Note that *MPLS* LSPs can still be used to transport packets to routes learned via BGP through the use of the `enable bgp mpls-next-hop` command.

When enabled, LDP uses routes learned via BGP when establishing LDP LSPs. As each established LSP consumes internal resources, it is recommended that this setting be used only in BGP environments where the number of BGP routes is controlled.

Example

The following command disables the use of BGP routes by LDP:

```
disable mpls ldp bgp-routes
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

disable mpls ldp loop-detection

```
disable mpls ldp loop-detection
```

Description

Disables LDP loop detection on the switch.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

Loop detection provides a mechanism for finding looping LSPs and for preventing Label Request messages from looping in the presence of non-merge-capable LSRs. The mechanism makes use of Path Vector and Hop Count TLVs carried by Label Request and Label Mapping messages.

When LDP loop detection is disabled, LDP does not attempt to detect routing loops.

Example

The following command globally disables LDP loop detection on the switch:

```
disable mpls ldp loop-detection
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

disable mpls ldp

```
disable mpls ldp [{vlan} vlan_name | vlan all]
```

Description

Disables LDP for the specified [MPLS](#)-configured [VLANs](#).

Syntax Description

vlan	Disables LDP for one or more specific VLANs.
<i>vlan_name</i>	Disables LDP on the specified VLAN.
vlan all	Disables LDP for all VLANs that have been added to MPLS.

Default

Disabled.

Usage Guidelines

When LDP is disabled, all LDP-advertised labels are withdrawn and all LDP peer sessions are terminated on the specified VLAN(s). By default, LDP is disabled for all VLANs. Specifying the optional all keyword disables LDP for all VLANs that have been added to MPLS.

Example

The following command disables LDP for all VLANs:

```
disable mpls ldp vlan all
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

disable mpls php

```
disable mpls php [{vlan} vlan_name | vlan all]
```

Description

Disables penultimate hop popping (PHP) on the specified *VLAN*. When enabled, PHP is requested on all LSPs advertised over that VLAN for which the switch is the egress LSR.

Syntax Description

vlan	Disables PHP for one or more specific VLANs.
<i>vlan_name</i>	Disables PHP on the specified VLAN.
vlan all	Disables PHP for all VLANs that have been added to <i>MPLS</i> .

Default

Disabled

Usage Guidelines

When PHP is disabled on a VLAN, penultimate hop popping is not requested on any LSPs advertised over that VLAN for which the switch is the egress LSR. Therefore, the Implicit Null Label is not used for any advertised mapping. Extreme's MPLS implementation always performs penultimate hop popping when requested to do so by a peer LSR. When the all VLANs option is selected, PHP is disabled on all existing MPLS interfaces.



Note

PHP is sometimes used to reduce the number of MPLS labels in use. If PHP is enabled on any MPLS interface, a unique MPLS label is consumed for every label advertised over that interface. Therefore, if PHP is being disabled to reduce label consumption, it should be done on all interfaces for minimal label consumption.

In ExtremeXOS, this command can be executed while MPLS is enabled.

Example

The following command disables penultimate hop popping (PHP) on the specified VLAN:

```
disable mpls php vlan vlan1
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

disable mpls protocol ldp

```
disable mpls protocol ldp
```

Description

Disables LDP for the switch.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

When LDP is disabled, all advertised LDP labels are withdrawn and LDP peer sessions are terminated. Note that this includes any LDP peer sessions established for L2 VPNs. By default, LDP is globally disabled. While LDP is transitioning to the enabled state, only the MPLS show commands are accepted.

Example

The following command globally disables LDP on the switch:

```
disable mpls protocol ldp
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

disable mpls protocol rsvp-te

```
disable mpls protocol rsvp-te
```

Description

Disables RSVP-TE for the switch.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

When RSVP-TE is disabled, all TE LSPs are released and TE LSPs cannot be established or accepted. While RSVP-TE is transitioning to the disabled state, only the [MPLS](#) show commands are accepted.

Example

The following command globally disables RSVP-TE on the switch:

```
disable mpls protocol rsvp-te
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

disable mpls rsvp-te bundle-message

```
disable mpls rsvp-te bundle-message [{vlan} vlan_name | vlan all]
```

Description

Disables the bundling of RSVP-TE messages for the specified [VLAN](#) interface.

Syntax Description

vlan	Specifies that message-bundling is to be disabled on a specific VLAN.
<i>vlan_name</i>	Identifies the VLAN interface on which message bundling is disabled.
vlan all	Specifies that message bundling is disabled on all VLAN interfaces that have been added to <i>MPLS</i> .

Default

Disabled.

Usage Guidelines

This command disables the bundling of RSVP-TE messages for the VLAN specified interface. By default, message bundling is disabled. Specifying the **all** keyword disables message bundling on all VLANs that have been added to MPLS.

Example

The following command disables message bundling on the specified VLAN:

```
disable mpls rsvp-te bundle-message vlan vlan_1
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

disable mpls rsvp-te fast-reroute

```
disable mpls rsvp-te fast-reroute
```

Description

Disables the *MPLS* RSVP-TE fast reroute (FRR) protection feature.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

When FRR is disabled on the LSR, all established FRR LSPs on the local LSR are torn down, and only standard LSPs can be signaled and processed. The configuration for any existing FRR LSPs is retained, but it is not used until the FRR protection feature is enabled. This command can be used to test the performance of an LSR without the FRR functionality or when the LSR doesn't behave as expected for either standard or FRR LSPs.

Example

The following command disables FRR protection on the local switch:

```
disable mpls rsvp-te fast-reroute
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

disable mpls rsvp-te lsp

```
disable mpls rsvp-te lsp [lsp_name | all]
```

Description

Disables an RSVP-TE LSP.

Syntax Description

<i>lsp_name</i>	Specifies the LSP within the switch to be disabled.
all	Disables all RSVP-TE configured LSPs.

Default

Enabled.

Usage Guidelines

This command disables an RSVP-TE LSP. When an RSVP-TE LSP is disabled, the switch terminates the LSP by signaling the destination by sending a PATH_TEAR message. If there are other LSPs configured to the same destination, traffic may continue to be transmitted to the destination over another LSP. Disabling an LSP does not otherwise change its configuration.

Example

The following command disables the LSP named lsp598:

```
disable mpls rsvp-te lsp lsp598
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

disable mpls rsvp-te summary-refresh

```
disable mpls rsvp-te summary-refresh [{vlan} vlan_name | vlan all]
```

Description

Disables the sending of summary refresh messages, instead of path messages, to refresh RSVP-TE path state for the specified [VLAN](#) interface.

Syntax Description

vlan	Specifies that summary refresh messages cannot refresh the RSVP-TE path state on one or more VLAN interfaces.
<i>vlan_name</i>	Specifies the VLAN interface for which RSVP-TE summary refresh messages are to be disabled.
vlan all	Specifies that summary refresh messages are to be disabled on all VLAN interfaces that have been added to MPLS .

Default

Disabled.

Usage Guidelines

This command disables the sending of summary refresh messages to refresh RSVP-TE path state for the specified VLAN interface. By default, summary refresh is disabled. Specifying the **all** keyword disables summary refresh on all VLANs that have been added to MPLS.

Example

The following command disables summary refresh on the specified VLAN:

```
disable mpls rsvp-te summary-refresh vlan vlan_1
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

disable mpls rsvp-te

```
disable mpls rsvp-te te [{vlan} vlan_name | vlan all]
```

Description

Disables RSVP-TE for the specified [MPLS](#)-configured [VLAN](#).

Syntax Description

vlan	Specifies that RSVP-TE is to be disabled on a specific VLAN.
<i>vlan_name</i>	Specifies the VLAN for which RSVP-TE is disabled.
vlan all	Disables RSVP-TE on all VLANs that have been added to MPLS.

Default

Disabled.

Usage Guidelines

This command disables RSVP-TE for the specified MPLS configured VLANs. When RSVP-TE is disabled, all TE LSPs are released and TE LSPs cannot be established or accepted. By default, RSVP-TE is disabled for all MPLS configured VLANs. Specifying the optional **all** keyword disables RSVP-TE for all VLANs that have been added to MPLS.

Example

The following command disables RSVP-TE on the named VLAN:

```
disable mpls rsvp-te vlan vlan_10
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

disable mpls static lsp

```
disable mpls static lsp {lsp_name | all }
```

Description

Administratively disables one or all static LSPs.

Syntax Description

<i>lsp_name</i>	Identifies an LSP to be disabled.
all	Specifies that all static LSPs on this LSR are to be disabled.

Default

N/A.

Usage Guidelines

On executing this command, the software de-activates the specified LSPs by setting the administrative state of each LSP to down.

Example

The following command disables a static LSP:

```
disable mpls static lsp lsp598
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

disable mpls vlan

```
disable mpls [{vlan} vlan_name | vlan all]
```

Description

Disables the *MPLS* interface for the specified *VLAN(s)*.

Syntax Description

vlan	Disables an MPLS interface for one or more specific VLANs.
<i>vlan_name</i>	Disables an MPLS interface on the specified VLAN.
vlan all	Disables an MPLS interface for all VLANs that have been added to MPLS.

Default

The MPLS interface is disabled for a VLAN.

Usage Guidelines

Disabling MPLS causes all LSPs to be released and all LDP and RSVP-TE peer sessions to be terminated on the specified VLAN(s).

Example

The following command disables an MPLS interface for the specified VLAN:

```
disable mpls vlan vlan-nyc
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

disable msdp

```
disable msdp {vr vrname}
```

Description

Disables *MSDP* on a virtual router.

Syntax Description

<i>vrname</i>	Specifies the name of the virtual router on which MSDP is being enabled or disabled. If a name is not specified, it is extracted from the current CLI context.
---------------	--

Default

MSDP is disabled by default.

Usage Guidelines

Use this command to disable MSDP on a virtual router.

Example

The following command disables MSDP on a virtual router:

```
disable msdp
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.ature-link-22.1"/>

disable msdp data-encapsulation

```
disable msdp data-encapsulation {vr vrname}
```

Description

Disables the encapsulation of locally originated SA messages with multicast data (if available).

Syntax Description

<i>vrname</i>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.
---------------	---

Default

By default, multicast data packet encapsulation is enabled for locally originated SA messages.

Usage Guidelines

None.

Example

The following command disables multicast data packet encapsulation:

```
disable msdp data-encapsulation
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the *MSDP* feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

disable msdp export local-sa

```
disable msdp export local-sa {vr vrname}
```

Description

Disables the advertisement of local sources to groups for which the router is an RP.

Syntax Description

<i>vrname</i>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.
---------------	---

Default

By default, the export of local sources is enabled. All sources are advertised if the router is an RP for the groups. Use this command to disable it.

Usage Guidelines

You can create a policy to filter out some of the local sources so that they are not advertised to *MSDP* peers and exposed to the external multicast domain. To configure an export filter, you must first disable the export of local sources (with the `disable msdp export local-sa` command), and then re-enable it with an export filter (with the `enable msdp export local-sa export-filter` command).

Example

The following example disables the advertisement of local sources:

```
disable msdp export local-sa
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

disable msdp peer

```
disable msdp [{peer} remoteaddr | peer all] {vr vr_name}
```

Description

Configures the administrative state of an *MSDP* peer.

Syntax Description

<i>remoteaddr</i>	Specifies the IP address of the MSDP peer to disable.
all	Disables all MSDP peers.
<i>vr_name</i>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

By default, MSDP peers are disabled.

Usage Guidelines

Use this command to administratively disable MSDP peers to stop exchanging SA messages.

Example

The following command disables an MSDP peer:

```
disable msdp peer 192.168.45.43
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

disable msdp process-sa-request

```
disable msdp [{peer} remoteaddr | peer all] process-sa-request {vr  
vrname}
```

Description

This command configures a router to reject SA request messages from a specified peer or all peers.

Syntax Description

peer all	Specifies all <i>MSDP</i> peers.
<i>remoteaddr</i>	Specifies the IP address of the MSDP peer.
<i>vrname</i>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

By default, all SA request messages are accepted from all peers.

Usage Guidelines

Use this command to configure the router to reject SA request messages from a specified peer or all peers.

You cannot change an SA request filter while SA request processing is enabled for an MSDP peer. You must first disable SA request processing for a peer and then re-enable it with an SA request filter.

You can use the following policy attributes in an SA request policy. All other attributes are ignored.

- Match:
 - multicast-group
 - multicast-source
 - pim-rp
- Set:
 - permit
 - deny

Example

The following example disables processing of SA request messages received from a peer with the IP address 192.168.45.43:

```
disable msdp peer 192.168.45.43 process-sa-request
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

disable msrp

```
disable msrp
```

Description

Disables MSRP on the switch.

Syntax Description

msrp	Multiple Stream Registration Protocol.
-------------	--

Default

Disabled.

Usage Guidelines

Use this command to disable MSRP on a switch.

Example

The following command disables MSRP:

```
disable msrp
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on all platforms if the AVB feature pack license is installed on the switch.

disable mvr

```
disable mvr
```

Description

Disables MVR on the system.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

None.

Example

The following example disables MVR on the system:

```
disable mvr
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MVR feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

disable mvrp

```
disable mvrp
```

Description

Disables MVRP globally on a switch.

Syntax Description

mvrp	Multiple VLAN Registration Protocol.
-------------	--

Default

Disabled.

Usage Guidelines

Use this command to disable MVRP globally on a switch. MVRP is run on the MVRP enabled ports only if the global setting is enabled. By default, MVRP is disabled globally and on individual ports. When MVRP is disabled globally, all MVRP packets are forwarded transparently.

Example

The following command disables MVRP:

```
disable mvrp
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable mvrp ports

```
disable mvrp ports [port_list | all]
```

Description

Disable MVRP on a given set of ports.

Syntax Description

mvrp	Multiple VLAN Registration Protocol.
<i>port_list</i>	Port(s) on which MVRP is to be enabled.
all	All ports.

Default

Disabled.

Usage Guidelines

Use this command to disable MVRP on given set of ports. MVRP is run on the MVRP enabled ports only if the global setting is enabled. By default MVRP is disabled globally and on individual ports. When MVRP is disabled globally, all MVRP packets will be forwarded transparently.

Example

The following command disables MVRP on ports 4 and 5:

```
disable mvrp ports 4-5
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable neighbor-discovery refresh

```
disable neighbor-discovery {vr vr_name} refresh
```

Description

Prevents the IPv6 neighbor cache from refreshing an entry before the timeout period expires.

Syntax Description

<i>vr_name</i>	Specifies a VR or VRF.
----------------	------------------------

Default

Enabled.

Usage Guidelines

None.

Example

The following example disables the refresh of neighbor discovery cache entries:

```
disable neighbor-discovery refresh
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

disable netlogin authentication failure vlan ports

```
disable netlogin authentication failure vlan ports [ports | all]
```

Description

Disables the configured authentication failure VLAN on the specified ports.

Syntax Description

all	Specifies all ports included in the authentication failure VLAN.
<i>ports</i>	Specifies one or more ports or slots and ports on which the authentication failure VLAN is enabled.

Default

All ports.

Usage Guidelines

Use this command to disable the configured authentication failure VLAN on either the specified ports, or all ports.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable netlogin authentication service-unavailable vlan ports

```
disable netlogin authentication service-unavailable vlan ports [ports |  
all]
```

Description

Disable the configured authentication service-unavailable VLAN on the specified ports.

Syntax Description

ports	Specifies one or more ports or slots and ports on which the authentication service-unavailable VLAN is enabled.
all	Specifies all ports included in the authentication service-unavailable VLAN.

Default

All ports.

Usage Guidelines

Use this command to disable the configured authentication service-unavailable VLAN on the specified ports, or on all ports.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable netlogin dot1x guest-vlan ports

```
disable netlogin dot1x guest-vlan ports [all | ports]
```

Description

Disables the guest VLAN on the specified 802.1X network login ports.

Syntax Description

all	Specifies all ports included in the guest VLAN.
<i>ports</i>	Specifies one or more ports included in the guest VLAN.

Default

Disabled.

Usage Guidelines

Use this command to disable the guest VLAN feature.

Enabling Guest VLANs

To enable the guest VLAN, use the following command:

```
enable netlogin dot1x guest-vlan ports [all | ports]
```

Example

The following command disables the guest VLAN on all ports:

```
disable netlogin dot1x guest-vlan ports all
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable netlogin logout-privilege

```
disable network login logout-privilege
```

Description

Disables network login logout window pop-up.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

This command controls the logout window pop-up on the web-based network client. This command applies only to the web-based authentication mode of network login. When disabled, the logout window pop-up will no longer appear. However, if session refresh is enabled, the login session will be terminated after the session refresh timeout.

Example

The following command disables network login logout-privilege:

```
disable netlogin logout-privilege
```


History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable netlogin ports

```
disable netlogin ports ports [{dot1x} {mac} {web-based}]
```

Description

Disables network login on a specified port for a particular method.

Syntax Description

<i>ports</i>	Specifies the ports for which network login should be disabled.
dot1x	Specifies 802.1X authentication.
mac	Specifies MAC-based authentication.
web-based	Specifies web-based authentication.

Default

Network login is disabled by default.

Usage Guidelines

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

This command applies to the MAC-based, web-based, and 802.1X mode of network login. To control which authentication mode is used by network login, use the following commands:

```
enable netlogin [{dot1x} {mac} {web-based}] disable netlogin [{dot1x} {mac} {web-based}]
```

Example

The following command disables dot1x and web-based network login on port 2:9:

```
disable netlogin ports 2:9 dot1x web-based
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable netlogin reauthenticate-on-refresh

```
disable netlogin reauthenticate-on-refresh
```

Description

Disables network login reauthentication on refresh.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

The web-based Netlogin client's session is periodically refreshed by sending an HTTP request which acts as a keep-alive without actually re-authenticating the user's credentials with the back-end *RADIUS* server or local database. If reauthenticate-on-refresh is enabled, re-authentication occurs with the session refresh.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable netlogin redirect-page

```
disable netlogin redirect-page
```

Description

Disables the network login redirect page function.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

This command disables the network login redirect page so that the client is sent to the originally requested page.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable netlogin session-refresh

```
disable netlogin session-refresh
```

Description

Disables network login session refresh.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

Network login sessions can refresh themselves after a configured timeout. After the user has been logged in successfully, a logout window opens which can be used to close the connection by clicking on the LogOut link. Any abnormal closing of this window is detected on the switch and the user is logged out after a time interval as configured for session refresh. The session refresh is enabled and set to three minutes by default.

This command applies only to the web-based authentication mode of network login.

Example

The following command disables network login session refresh:

```
disable netlogin session-refresh
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable netlogin

```
disable netlogin [{dot1x} {mac} {web-based}]
```

Description

Disables network login modes.

Syntax Description

dot1x	Specifies 802.1X authentication.
mac	Specifies MAC-based authentication.
web-based	Specifies web-based authentication.

Default

All types of authentication are disabled.

Usage Guidelines

Any combination of authentication types can be disabled on the same switch. To enable an authentication mode, use the following command:

```
enable netlogin [{dot1x} {mac} {web-based}]
```

Example

The following command disables MAC-based network login:

```
disable netlogin mac
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable network-clock gtp ports

```
disable network-clock gtp ports [port_list {only} | all]
```

Description

Disables gPTP on one or more ports.

Syntax Description

<i>port_list</i>	Specifies one or more the the switch's physical ports.
only	Apply change only to specified port, even if port is master of a load sharing group.
all	Specifies all of the switch's physical ports.

Default

Disabled.

Usage Guidelines

Use this command to configure on which ports gPTP runs. gPTP runs on no ports if it is not enabled in the switch by `enable network-clock gtp`.

Example

```
disable network-clock gtp ports 1-3
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on all platforms if the AVB feature pack license is installed on the switch.

disable network-clock gtp

```
disable network-clock gtp
```

Description

Disables gPTP on the switch.

Syntax Description

network-clock	Network clock.
gptp	IEEE 802.1AS Generalized Precision Time Protocol (gPTP).

Default

Disabled.

Usage Guidelines

Use this command to disable gPTP after having enabled it.

Example

```
disable network-clock gptp
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the AVB feature pack license and have it installed. For complete information about software licensing, including how to obtain and upgrade your license, and which platforms support the AVB feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

disable nodealias ports

```
disable nodealias ports [port_list | all]
```

Description

This command disables the Node Alias feature on specified ports. Node Alias discovers information about the end systems on a per-port basis. Information from packets from end systems, such as VLANID, source MAC address, source IP address, protocol, etc. are captured in a database that can be queried.

Syntax Description

nodealias	Node Alias feature that maps source IP address, MAC address, host name, and protocol on a per port basis.
ports	Designates that Node Alias should be disabled on specified ports.
<i>port_list</i>	Specifies on which ports to have Node Alias disabled. Designated as a port list separated by comma (,) or dash (-).
all	Specifies that all ports have Node Alias disabled.

Default

Node Alias is disabled by default on all ports.

Usage Guidelines

If the port is part of a [LAG](#), Node Alias should be disabled separately on each LAG port.

Example

The following example disables Node Alias on all ports:

```
disable nodealias ports all
```

History

This command was first available in ExtremeXOS 22.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable nodealias protocol

disable **nodealias protocol** [*protocol_name* | **any**]

Description

This command designates the specific protocols to remove from the list of detected protocols for the Node Alias feature. Node Alias discovers information about the end systems on a per-port basis. Information from packets from end systems, such as VLANID, source MAC address, source IP address, protocol, etc. are captured in a database that can be queried.

Syntax Description

nodealias	Node Alias feature that maps source IP address, MAC address, host name, and protocol on a per port basis.
protocol	Designates selection of protocols to detect.
<i>protocol_name</i>	Specifies disabling a protocol to detect (one at a time). The following protocols are enabled by default: IPv4, IPv6, <i>OSPF</i> , <i>BGP</i> , <i>VRRP</i> , DHCPs, DHCPc, BOOTPS, BOOTPC, UDP, BPDU, LLMNR, SSDP, and mDNS.
any	Specifies disabling all protocols.

Default

The following protocols are enabled by default: IPv4, IPv6, OSPF, BGP, VRRP, DHCP, DHCP, BOOTP, BOOTP, UDP, BPDU, LLNMR, SSDP, and mDNS.



Note

- ARP is categorized under IP.
- UDP entry is created when destination IP address is broadcast.
- BPDU means *STP* and GVRP frames.

Usage Guidelines

By default, the following protocols are enabled (IPv4, IPv6, OSPF, BGP, VRRP, DHCP, DHCP, BOOTP, BOOTP, UDP, BPDU, LLNMR, SSDP, mDNS). You can optionally disable any of these protocols (and then enable them back if desired).

Example

The following example disables BGP from being detected:

```
disable nodealias protocol bgp
```

History

This command was first available in ExtremeXOS 22.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable ntp

```
disable ntp
```

Description

Disables NTP globally on the switch.

Syntax Description

N/A.

Default

NTP is disabled by default.

Usage Guidelines

N/A.

Example

The following command disables NTP globally on the switch:

```
disable ntp
```

History

This command was first available in ExtremeXOS 12.7.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable ntp authentication

```
disable ntp authentication
```

Description

Disables NTP authentication globally on the switch.

Syntax Description

N/A.

Default

NTP authentication is disabled by default.

Usage Guidelines

If authentication is disabled, NTP will not use any authentication mechanism to a server or from clients. To use authentication for a specific server, enable NTP authentication globally, and then configure an RSA Data Security, Inc. *MD5* Message-Digest Algorithm or SHA256 key index for the specific server.

Example

The following command disables NTP authentication globally on the switch:

```
# disable ntp authentication
```

History

This command was first available in ExtremeXOS 12.7.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable ntp broadcast-client

```
disable ntp broadcast-client {{vr} vr_name}
```

Description

Disables an NTP broadcast client on the switch.

Syntax Description

broadcast-client	Specifies enabling NTP broadcast client.
vr	Specifies disabling NTP broadcast client for a VR.
<i>vr_name</i>	Specifies the VR name. If a VR name is not specified, the VR of current command context is used.

Default

An NTP broadcast client is enabled by default.

If a VR name is not specified, the VR of current command context is used.

Usage Guidelines

If the broadcast client function is enabled, the system can receive broadcast-based NTP messages and process them only if a VLAN is enabled for NTP and the VLAN is active.

Example

The following command disables an NTP broadcast client on the switch:

```
disable ntp broadcast client
```

History

This command was first available in ExtremeXOS 12.7.

The **vr** was added in ExtremeXOS 22.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable ntp broadcast-server

```
disable ntp {vlan} vlan-name broadcast-server
```

Description

Prevents NTP from sending broadcast messages to a VLAN.

Syntax Description

<i>vlan-name</i>	Specifies the name of a particular VLAN.
------------------	--

Default

NTP does not send broadcast messages to a VLAN by default.

Usage Guidelines

N/A.

Example

The following command prevents NTP from sending broadcast messages to a VLAN called “Northwest”:

```
disable ntp vlan Northwest broadcast-server
```

History

This command was first available in ExtremeXOS 12.7.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable ntp vlan

```
disable ntp [{vlan} vlan-name | all] [{vr} vr_name]
```

Description

Disables NTP on a VLAN.

Syntax Description

disable	Disables NTP on a VLAN.
<i>vlan-name</i>	Specifies the name of a particular VLAN on which to enable or disable NTP.
all	Enables or disables NTP on all VLANs.
vr	Specifies disabling NTP on a VR.
<i>vr_name</i>	Specifies the VR name to disable NTP on. If a VR name is not specified, the VR of current command context is used.

Default

NTP is disabled on all VLANs by default.

Usage Guidelines

N/A.

Example

The following command disables NTP on all VLANs:

```
disable ntp all
```

The following command disables NTP on specific VLAN:

```
disable ntp vlan vlan-1
```

History

This command was first available in ExtremeXOS 12.7.

The **vr** option was added in ExtremeXOS 22.2

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable ntp vr

```
disable ntp vr vr_name
```

Description

This command disables NTP from the specified VR.

Syntax Description

vr	Specifies disabling NTP on a VR.
<i>vr_name</i>	Specifies the VR name to disable NTP from. If a VR name is not specified, the VR of current command context is used.

Default

If a VR name is not specified, the VR of current command context is used.

Example

The following example disables NTP from a VR named "vr1".

```
disable ntp vr vr1
```

History

This command was first available in ExtremeXOS 22.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable ospf

```
disable ospf
```

Description

Disables the *OSPF* process for the router.

Syntax Description

This command has no keywords or arguments.

Default

N/A.

Usage Guidelines

Not applicable.

Example

The following command disables the OSPF process for the router:

```
disable ospf
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [***Switch Engine 32.3 Feature License Requirements***](#) document.

disable ospf capability opaque-lsa

```
disable ospf capability opaque-lsa
```

Description

Disables opaque LSAs across the entire system.

Syntax Description

This command has no keywords or arguments.

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 first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

disable ospf export

```
disable ospf export [bgp | direct | host-mobility | e-bgp | i-bgp | rip |
static | isis | isis-level-1 | isis-level-1-external | isis-level-2 |
isis-level-2-external]
```

Description

Disables redistribution of routes to [OSPF](#).

Syntax Description

bgp	Specifies BGP routes.
direct	Specifies direct routes.
host-mobility	Specifies host mobility.
e-bgp	Specifies E-BGP routes.
i-bgp	Specifies I-BGP routes.
rip	Specifies RIP routes.
static	Specifies static routes.
isis	Specifies IS-IS routes.
isis-level-1	Specifies ISIS Level 1 routes.
isis-level-1-external	Specifies ISIS Level 1 External routes.
isis-level-2	Specifies ISIS Level 2 routes.
isis-level-2-external	Specifies ISIS Level 2 External routes.

Default

The default setting is disabled.

Usage Guidelines

Use this command to stop OSPF from exporting routes derived from other protocols.

Example

The following command disables OSPF to export BGP-related routes to other OSPF routers:

```
disable ospf export bgp
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [**Switch Engine 32.3 Feature License Requirements**](#) document.

NEW! disable ospf export vr

```
disable ospf export {vr} vr-name route-type
```

Description

Disables redistribution of routes between [*OSPF*](#) instances.

Syntax Description

vr	Specifies Virtual Router.
<i>vr-name</i>	Specifies the source Virtual Router for exported routes.
<i>route-type</i>	Specifies the OSPF inter-VR export route type.

Default

The default setting is disabled.

Usage Guidelines

Use this command to stop OSPF from exporting routes between other protocols.

Example

The following command disables OSPF to redistribute routes between OSPF instances:

```
disable ospf export vr ??
```

History

This command was first available in ExtremeXOS 32.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable ospf mpls-next-hop

```
disable ospf mpls-next-hop {vr vrf_name}
```

Description

Disables IP forwarding over calculated *MPLS* LSPs to subnets learned via *OSPF*.

Syntax Description

<i>vrf_name</i>	Specifies OSPF on a particular VRF.
-----------------	-------------------------------------

Default

Disabled.

Usage Guidelines

This command disables IP forwarding over calculated MPLS LSPs to subnets learned via OSPF. (Calculated refers to an LSP that only reaches part of the way to the destination). By default, IP forwarding over MPLS LSPs to subnets learned via OSPF is disabled.

In order to disable OSPF on a particular VRF, you must supply the optional *vr vrf_name* CLI parameter.

Example

The following command disables OSPF's use of MPLS LSPs to reach OSPF routes:

```
disable ospf mpls-next-hop
```

History

This command was first available in ExtremeXOS 11.6.

The **vr** keyword and *vrf_name* variable were added in ExtremeXOS 15.3.

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

disable ospf originate-default

```
disable ospf originate-default
```

Syntax Description

There are no keywords or variables for this command.

Default

Not applicable.

Usage Guidelines

Not applicable.

Example

The following command disables generating a default external LSA:

```
disable ospf originate-default
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

disable ospf restart-helper-lsa-check

```
disable ospf [vlan [all | vlan-name] | area area-identifier | virtual-link router-identifier area-identifier] restart-helper-lsa-check
```

Description

Disables the restart helper router from terminating graceful [OSPF](#) restart when received LSAs would affect the restarting router.

Syntax Description

all	Specifies all <u>VLANs</u> .
<i>vlan-name</i>	Specifies a VLAN name.
<i>router-identifier</i>	Specifies the router ID of the remote router of the virtual link.
<i>area-identifier</i>	Specifies an OSPF area.

Default

The default is enabled.

Usage Guidelines

This command disables the restart helper router from terminating graceful OSPF restart when received LSAs would affect the restarting router.

Example

The following command disables a router from terminating graceful OSPF restart for all routers in area 10.20.30.40 if it receives an LSA that would affect routing:

```
disable ospf area 10.20.30.40 restart-helper-lsa-check
```

History

This command was first available in ExtremeXOS 11.3.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the ***Switch Engine 32.3 Feature License Requirements*** document.

disable ospf use-ip-router-alert

```
disable ospf use-ip-router-alert
```

Description

Disables the router alert IP option in outgoing OSPF control packets.

Syntax Description

This command has no keywords or arguments.

Default

Disabled.

Usage Guidelines

Not applicable.

Example

The following command disables the OSPF router alert IP option:

```
disable ospf use-ip-router-alert
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [***Switch Engine 32.3 Feature License Requirements***](#) document.

disable ospf vxlan-extensions

```
disable ospf vxlan-extensions
```

Description

This command disables the OSPFv2 VXLAN extensions.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

N/A.

Example

```
# disable ospf vxlan-extensions
```

History

This command was first available in ExtremeXOS 21.1.

Platform Availability

This command is supported on ExtremeSwitching 5420 and 5520 series switches, and stacks with 5420 and 5520 slots only.

disable ospfv3

```
disable ospfv3
```

Description

Disables *OSPFv3* for the router.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command disables OSPFv3 for the router:

```
disable ospfv3
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Advanced Edge or Core license as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

disable ospfv3 restart-helper-lsa-check

```
disable ospfv3 [[vlan | tunnel] all | vlan vlan-name | {tunnel} tunnel-name | area area-identifier] restart-helper-lsa-check
```

Description

This command configures the restart helper router to terminate *OSPFv3* graceful restart when received LSAs would affect the restarting router. This occurs when the restart helper receives an LSA that is flooded to the restarting router or when there is a changed LSA on the restarting router's retransmission list when graceful restart is initiated.

Syntax Description

vlan	<i>VLAN</i> .
all	All VLANs.
<i>vlan-name</i>	VLAN name.
area	OSPFv3 area.
<i>area-identifier</i>	Area identifier.
restart-helper-lsa-check	Terminate graceful restart mode when there is a change to an LSA.

Default

LSA check is enabled by default.

History

This command was first available in ExtremeXOS 21.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable ospfv3 export

```
disable ospfv3 export [direct | ripng | static | isis | isis-level-1
| isis-level-1-external | isis-level-2 | isis-level-2-external | bgp
e-bgp i-bgp]
```

Description

Disables redistribution of routes to *OSPFv3*.

Syntax Description

direct	Specifies direct routes.
ripng	Specifies <i>RIP</i> routes.
static	Specifies static routes.
isis	Specifies IS-IS routes.
isis-level-1	Specifies IS-IS Level 1 routes.

isis-level-1-external	Specifies IS-IS Level 1 External routes.
isis-level-2	Specifies IS-IS Level 2 routes.
isis-level-2-external	Specifies IS-IS Level 2 External routes.
bgp	Specifies <i>BGP</i> IPv6 routes.
e-bgp	Specifies EBGP routes.
i-bgp	Specifies EBGP routes.

Default

The default setting is disabled.

Usage Guidelines

Use this command to stop OSPFv3 from exporting routes derived from other protocols.

Example

The following command disables OSPFv3 to export *RIPng* routes to other OSPFv3 routers:

```
disable ospfv3 export ripng
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the *Switch Engine 32.3 Feature License Requirements* document.

disable ospfv3 virtual-link restart-helper-lsa-check

```
disable ospfv3 virtual-link {routerid} router-identifier {area} area-identifier restart-helper-lsa-check
```

Description

This command configures the restart helper router to terminate *OSPF* graceful restart when received LSAs would affect the restarting router. This occurs when the restart helper receives an LSA that will be flooded to the restarting router or when there is a changed LSA on the restarting router's retransmission list when graceful restart is initiated.

Syntax Description

virtual-link	<i>OSPFv3</i> virtual link.
routerid	OSPFv3 router ID.
<i>router-identifier</i>	Router ID of neighbor OSPFv3 router.
area	OSPFv3 area.
<i>area-identifier</i>	Transit area ID of virtual link.
restart-helper-lsa-check	Terminates graceful restart helper mode when there is a change to an LSA (default is enabled).

Default

Enabled.

History

This command was first available in ExtremeXOS 21.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable pim iproute sharing

```
disable pim {ipv4 | ipv6} iproute sharing
```

Description

Disables the PIM *ECMP* feature.

Syntax Description

iproute	IP Route
sharing	Equal Cost Multipath Routing

Default

Disabled.

Usage Guidelines

None.

Example

The following command disables the PIM ECMP feature:

```
disable pim ipv4 iproute sharing
```

History

This command was first available in ExtremeXOS 15.3.2.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

disable pim snooping

```
disable pim snooping {{vlan} name}
```

Description

Disables PIM snooping and clears all the snooping PIM neighbors, joins received on the VLAN, and the forwarding entries belonging to one or all VLANs.

Syntax Description

<i>name</i>	Specifies a VLAN.
-------------	-------------------

Default

Disabled.

Usage Guidelines

None.

Example

The following command disables PIM snooping for all VLANs on the switch:

```
disable pim snooping
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

NEW! disable pim ssm tunnel

```
disable pim {ipv4} ssm tunnel [tunnel_name | tunnel all]
```

Description

Disables PIM SSM tunnels on a router interface.

Syntax Description

ipv4	Specifies the IPv4 address family.
ssm	Specifies PIM SSM mode.
tunnel	Specifies the tunnel name.
<i>tunnel_name</i>	Specifies to configure PIM information for interfaces.
all	Specifies to configure PIM information on all Layer 3 interfaces.

Default

Disabled on all interfaces.

Usage Guidelines

This command disables PIM-SSM on the specified Layer 3 VLAN.

IGMPv3 include messages for multicast addresses in the SSM range is only processed by PIM if PIM-SSM is enabled on the interface. Any non-IGMPv3 messages in the SSM range are not processed by PIM on any switch interface, whether SSM is enabled or not.

Example

The following example disables PIM-SSM multicast routing on tunnel accounting:

```
disable pim ssm tunnel accounting
```

History

This command was first available in ExtremeXOS 32.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

disable pim ssm vlan

```
disable pim {ipv4 | ipv6} ssm vlan [vlan_name | all]
```

Description

Disables PIM SSM on a router interface.

Syntax Description

ipv4	Specifies the IPv4 address family.
ipv6	Specifies the IPv6 address family.
<i>vlan_name</i>	Specifies a <u>VLAN</u> name.
all	Specifies all VLANs.

Default

Disabled on all interfaces.

Usage Guidelines

This command disables PIM-SSM on the specified Layer 3 VLAN.

IGMPv3 include messages for multicast addresses in the SSM range is only processed by PIM if PIM-SSM is enabled on the interface. Any non-IGMPv3 messages in the SSM range are not processed by PIM on any switch interface, whether SSM is enabled or not.

Example

The following example disables PIM-SSM multicast routing on VLAN accounting:

```
disable pim ssm vlan accounting
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

disable pim

```
disable pim {ipv4 | ipv6}
```

Description

Disables PIM on the system.

Syntax Description

ipv4	Specifies the IPv4 address family.
ipv6	Specifies the IPv6 address family.

Default

Disabled.

Usage Guidelines

None.

Example

The following example disables PIM on the system:

```
disable pim ipv4
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

disable policy

```
disable policy
```

Description

This command disables the ONEPolicy functionality.

Syntax Description

This command has no arguments or variables.

Default

None.

Usage Guidelines

None.

Example

The following example shows how to disable ONEPolicy:

```
x450G2-48t-10G4.4 # disable policy
```

History

This command was first available in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable port

disable port [*port_list* | **all**]

Description

Disables one or more ports on the switch.

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
all	Specifies all ports on the switch.

Default

Enabled.

Usage Guidelines

Use this command for security, administration, and troubleshooting purposes.

When a port is disabled, the link is brought down.

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 ports 3, 5, and 12 through 15 on a switch:

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

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable ports mlag-id

```
disable ports [mlag-id mlag_id]
```

Description

Disables the current ports associated with the given ID.

Syntax Description

mlag-id	Port associated with MLAG.
<i>mlag_id</i>	MLAG identifier value of the MLAG port. Range is 1–65,000.

Default

N/A.

Usage Guidelines

If any ports are added or deleted from the LAG, the port state for those ports is not changed.

In MLAG orchestration mode, this command is executed on the other MLAG peer before it is executed on the MLAG peer on which the command is run. In orchestration mode, if the MLAG port numbers are not same on both the peers, it is possible that a different set of port numbers are disabled on the different MLAG peers. This command helps ensure that the correct set of ports associated with the MLAG ID are disabled.

If the port associated with the given MLAG ID is a load shared port, all the member ports associated with this load shared group are disabled.

If the port associated with the given MLAG ID is a virtual port, the command is ignored.

Example

The following example disables the ports associated with MLAG ID "123":

```
# disable ports mlag-id 123
```

History

This command was first available in ExtremeXOS 32.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, 5720 series switches.

disable radius

```
disable radius {mgmt-access | netlogin}
```

Description

Disables the *RADIUS* client.

Syntax Description

mgmt-access	Specifies the switch management RADIUS authentication server.
netlogin	Specifies the network login RADIUS authentication server.

Default

RADIUS authentication is disabled for both switch management and network login by default.

Usage Guidelines

Use the mgmt-access keyword to disable RADIUS authentication for switch management functions.

Use the netlogin keyword to disable RADIUS authentication for network login.

If you do not specify a keyword, RADIUS authentication is disabled on the switch for both management and network login.

Example

The following command disables RADIUS authentication on the switch for both management and network login:

```
disable radius
```

The following command disables RADIUS authentication on the switch for network login:

```
disable radius netlogin
```

History

This command was first available in ExtremeXOS 10.1.

The mgmt-access and netlogin keywords were added in ExtremeXOS 11.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable radius-accounting

```
disable radius-accounting {mgmt-access | netlogin}
```

Description

Disables *RADIUS* accounting.

Syntax Description

mgmt-access	Specifies the switch management RADIUS accounting server.
netlogin	Specifies the network login RADIUS accounting server.

Default

RADIUS accounting is disabled for both switch management and network login by default.

Usage Guidelines

Use the mgmt-access keyword to disable RADIUS accounting for switch management functions.

Use the netlogin keyword to disable RADIUS accounting for network login.

If you do not specify a keyword, RADIUS accounting is disabled on the switch for both management and network login.

Example

The following command disables RADIUS accounting on the switch for both management and network login:

```
disable radius-accounting
```


The following command disables RADIUS accounting on the switch for network login:

```
disable radius-accounting netlogin
```

History

This command was first available in ExtremeXOS 10.1.

The mgmt-access and netlogin keywords were added in ExtremeXOS 11.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable radius dynamic-authorization

```
disable radius dynamic-authorization
```

Description

Disables dynamic authorization on *RADIUS* client.

Syntax Description

This command has no arguments or variables.

Default

RADIUS dynamic authorization is disabled by default.

Example

The following command disables dynamic authorization RADIUS authentication on the switch:

```
disable radius dynamic-authorization
```

History

This command was first available in ExtremeXOS 22.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

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 first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms with a Base license or higher as described in the [***Switch Engine 32.3 Feature License Requirements***](#) document.

disable rip aggregation

```
disable rip aggregation
```

Description

Disables the *RIP* aggregation of subnet information on a RIP version 2 (RIPv2) router.

Syntax Description

This command has no arguments or variables.

Default

RIP aggregation is disabled by default.

Usage Guidelines

The disable RIP aggregation command disables the RIP aggregation of subnet information on a switch configured to send RIPv2-compatible traffic. The switch summarizes subnet routes to the nearest class network route. The following rules apply when using RIP aggregation:

- Within a class boundary, no routes are aggregated.
- 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 first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms with a Base license or higher as described in the [*Switch Engine 32.3 Feature License Requirements*](#) document.

disable rip export

```
disable rip export [ bgp | direct | e-bgp | i-bgp | ospf | ospf-extern1 | ospf-extern2 | ospf-inter | ospf-intra | static | isis | isis-level-1 | isis-level-1-external | isis-level-2 | isis-level-2-external ]
```

Description

Disables *RIP* from redistributing routes from other routing protocols.

Syntax Description

bgp	Specifies <i>BGP</i> routes.
direct	Specifies interface routes (only interfaces that have IP forwarding enabled are exported).
e-bgp	Specifies external <i>BGP</i> routes.
i-bgp	Specifies internal BGP routes.
ospf	Specifies all <i>OSPF</i> routes.
ospf-extern1	Specifies OSPF external route type 1.

ospf-extern2	Specifies OSPF external route type 2.
ospf-inter	Specifies OSPF-inter area routes.
ospf-intra	Specifies OSPF-intra area routes.
static	Specifies static routes.
isis	Specifies IS-IS routes.
isis-level-1	Specifies ISIS Level 1 routes.
isis-level-1-external	Specifies ISIS Level 1 External routes.
isis-level-2	Specifies ISIS Level 2 routes.
isis-level-2-external	Specifies ISIS Level 2 External routes.

Default

Disabled.

Usage Guidelines

This command disables the exporting of BGP, static, direct, and OSPF-learned routes into the RIP domain.

Example

The following command disables RIP from redistributing any routes learned from OSPF:

```
# disable rip export ospf
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms with a Base license or higher as described in the [***Switch Engine 32.3 Feature License Requirements***](#) document.

disable rip originate-default

```
disable rip originate-default
```

Description

Disables the advertisement of a default route.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

None.

Example

The following command unconfigures a default route to be advertised by RIP if no other default route is advertised:

```
# disable rip originate-default
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms with a Base license or higher as described in the **Switch Engine 32.3 Feature License Requirements** document.

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 first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms with a Base license or higher as described in the [***Switch Engine 32.3 Feature License Requirements***](#) document.

disable rip splithorizon

```
disable rip splithorizon
```

Description

Disables the split horizon algorithm for *RIP*.

Syntax Description

This command has no arguments or variable.

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 first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms with a Base license or higher as described in the [**Switch Engine 32.3 Feature License Requirements**](#) document.

disable rip triggerupdates

```
disable rip triggerupdates
```

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 first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms with a Base license or higher as described in the [**Switch Engine 32.3 Feature License Requirements**](#) document.

disable rip use-ip-router-alert

```
disable rip use-ip-router-alert
```

Description

Disables router alert IP option in outgoing *RIP* control packets.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

None.

Example

The following command disables the RIP router alert IP option:

```
# disable rip use-ip-router-alert
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms with a Base license or higher as described in the [***Switch Engine 32.3 Feature License Requirements***](#) document.

disable ripng

```
disable ripng
```

Description

Disables *RIPng* for the whole router.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

None.

Example

The following command disables RIPng for the whole router:

```
disable ripng
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Edge, Advanced Edge, or Core license. For licensing information, see the [Switch Engine 32.3 Feature License Requirements](#) document.

disable ripng export

```
disable ripng export [direct | ospfv3 | ospfv3-extern1 | ospfv3-extern2
| ospfv3-inter | ospfv3-intra | static | isis | isis-level-1 | isis-
level-1-external | isis-level-2 | isis-level-2-external | bgp e-bgp i-
bgp]
```

Description

Disables *RIPng* from redistributing routes from other routing protocols.

Syntax Description

direct	Specifies directly reachable subnets from the router (only interfaces that have IP forwarding enabled are exported).
ospfv3	Specifies all <i>OSPFv3</i> routes.
ospfv3-extern1	Specifies OSPFv3 external route type 1.
ospfv3-extern2	Specifies OSPFv3 external route type 2.
ospfv3-inter	Specifies OSPFv3-inter area routes.
ospfv3-intra	Specifies OSPFv3-intra area routes.
static	Specifies user configured static routes.
isis	Specifies IS-IS routes.
isis-level-1	Specifies IS-IS Level 1 routes.
isis-level-1-external	Specifies IS-IS Level 1 External routes.
isis-level-2	Specifies IS-IS Level 2 routes.
isis-level-2-external	Specifies IS-IS Level 2 External routes.

bgp	Specifies <i>BGP</i> IPv6 routes
e-bgp	Specifies EBGP routes.
i-bgp	Specifies IBGP routes.

Default

Disabled.

Usage Guidelines

This command disables the exporting of static, direct, IS-IS, and *OSPF*-learned routes from the switch routing table into the RIPng domain.

Example

The following command disables RIPng from redistributing any routes learned from OSPFv3:

```
disable ripng export ospfv3
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Edge, Advanced Edge, or Core license. For licensing information, see the [Switch Engine 32.3 Feature License Requirements](#) document.

disable ripng originate-default

```
disable ripng originate-default
```

Description

Disables the advertisement of a default route to the neighbors.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

None.

Example

The following command unconfigures a default route to be advertised by *RIPng* if no other default route is advertised:

```
disable ripng originate-default
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Edge, Advanced Edge, or Core license. For licensing information, see the [Switch Engine 32.3 Feature License Requirements](#) document.

disable ripng poisonreverse

```
disable ripng poisonreverse {vlan vlan-name | tunnel tunnel_name | [vlan  
| tunnel] all}
```

Description

Disables poison reverse algorithm for *RIPng*.

Syntax Description

<i>vlan-name</i>	Specifies an IPv6 configured <i>VLAN</i> .
<i>tunnel_name</i>	Specifies an IPv6 tunnel.
all	Specifies all interfaces.

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 RIPng:

```
disable ripng poisonreverse
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Edge, Advanced Edge, or Core license. For licensing information, see the [Switch Engine 32.3 Feature License Requirements](#) document.

disable ripng splithorizon

```
disable ripng splithorizon {vlan vlan-name | tunnel tunnel_name | [vlan  
| tunnel] all}
```

Description

Disables the split horizon algorithm for *RIPng*.

Syntax Description

<i>vlan-name</i>	Specifies an IPv6 configured <i>VLAN</i> .
<i>tunnel_name</i>	Specifies an IPv6 tunnel.
all	Specifies all interfaces.

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 RIPng:

```
disable rip splithorizon
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Edge, Advanced Edge, or Core license. For licensing information, see the [Switch Engine 32.3 Feature License Requirements](#) document.

disable ripng triggerupdate

```
disable ripng triggerupdate {vlan vlan-name | tunnel tunnel_name | [vlan
| tunnel] all}
```

Description

Triggered updates are a mechanism for immediately notifying a router's neighbors when the router adds or deletes routes or changes their metric. This command disables the trigger update mechanism.

Syntax Description

<i>vlan-name</i>	Specifies an IPv6 configured <u>VLAN</u> .
<i>tunnel_name</i>	Specifies an IPv6 tunnel.
all	Specifies all interfaces.

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 RIPng-related traffic.

When this feature is disabled, any metric change on the interface, or an interface going down will not be communicated until the next periodic update. To configure how often periodic updates are sent, use the following command:

```
configure ripng updatetime
```

Example

The following command disables the trigger update mechanism:

```
disable ripng triggerupdate
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Edge, Advanced Edge, or Core license. For licensing information, see the [Switch Engine 32.3 Feature License Requirements](#) document.

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 queries of statistics. Collecting of history, alarms, and events is stopped; however, the switch still queries old data.

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.

To view the RMON memory usage statistics for a specific memory type (for example, statistics, events, logs, history, or alarms) or for all memory types, use the following command:

```
show rmon memory {detail | memoryType}
```

Example

The following command disables the collection of RMON statistics on the switch:

```
disable rmon
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable router-discovery

```
disable router-discovery {ipv6} vlan vlan_name
```

Description

Disables router discovery advertisements on the VLAN and the processing of router discovery messages.

Syntax Description

<i>vlan_name</i>	Specifies an IPv6 configured VLAN.
------------------	------------------------------------

Default

N/A.

Usage Guidelines

None.

Example

The following example disables router discovery for the VLAN "top_floor":

```
disable router-discovery vlan top_floor
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

disable sflow ports

```
disable sflow ports port_list
```

Description

Disables sFlow statistical packet sampling and statistics gathering on a particular list of ports.

Syntax Description

<code>port_list</code>	Specifies a list of ports.
------------------------	----------------------------

Default

Disabled.

Usage Guidelines

This command disables sFlow on a particular list of ports. Once sFlow is disabled on a port, sampling and polling will stop. If sFlow is disabled globally, all sampling and polling stops.

Use the following command to disable sFlow globally:

```
disable sflow
```

Example

The following command disables sFlow sampling on port 3:1:

```
disable sflow ports 3:1
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable sflow

```
disable sflow
```

Description

Globally disables sFlow statistical packet sampling.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command disables sFlow globally on the switch. When you disable sFlow globally, the individual ports are also put into the disabled state. If you later enable the global sFlow state, individual ports return to their previous state.

Example

The following command disables sFlow sampling globally:

```
disable sflow
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable sharing

disable sharing *port*

Description

Disables a load-sharing group of ports, also known as a LAG.

Syntax Description

<i>port</i>	Specifies the logical port of a load-sharing group or link aggregation group (LAG). Specifies a port or a combination of the slot and port number.
-------------	--

Default

Disabled.

Usage Guidelines

When sharing is disabled, the logical port retains all configuration including VLAN membership. All other member ports are removed from all VLANs to prevent loops and their configuration is reset to default values.

Any attempt to disable sharing on ports that have configuration is denied with the following error message:

ERROR: Sharing configuration on MLAG ports cannot be modified. Use "disable mlag port <port>" to remove port from MLAG group first.

Example

The following command disables sharing on master logical port 9, which contains ports 9 through 12, on a switch:

```
disable sharing 9
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable slpp guard

```
disable slpp guard ports [port_list | all]
```

Description

Disable the Simple Loop Protection Protocol (SLPP) Guard feature.

Syntax Description

slpp	Specifies disabling SLPP.
guard	Specifies not using the SLPP Guard feature, which disables a port as soon as an SLPP PDU is received.
ports	Specifies selecting ports on which to disable SLPP guard.
<i>port_list</i>	Selects which ports on which to disable SLPP guard.
all	Specifies disabling SLPP guard on all ports.

Default

By default, SLPP Guard is disabled on all ports.

Usage Guidelines

SLPP is an application that detects loops in a Split Multi-link Trunking (SMLT) network. SLPP Guard is a complementary feature that helps prevent loops in networks by administratively disabling an edge port if a switch receive an SLPP PDU from an SMLT network.

Example

The following example disables SLPP Guard on port 5:

```
# disable slpp guard ports 5
```

History

This command was first available in ExtremeXOS 30.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable smartredundancy

```
disable smartredundancy port_list
```

Description

Disables the Smart Redundancy feature.

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
------------------	---

Default

Enabled.

Usage Guidelines

The Smart Redundancy feature works in concert with the software-controlled redundant feature. When Smart Redundancy is disabled, the switch attempts only to reset the primary port to active if the redundant port fails. That is, if you disable Smart Redundancy, the traffic does not automatically return to the primary port once it becomes active again; the traffic continues to flow through the redundant port even after the primary port comes up again.

Example

The following command disables the Smart Redundancy feature on ports 1 through 4 on a switch:

```
disable smartredundancy 1-4
```

History

This command was available in ExtremeXOS 11.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable snmp access vr

```
disable snmp access vr [vr_name | all]
```

Description

Selectively disables *SNMP* access on virtual routers.

Syntax Description

<i>vr_name</i>	Specifies the virtual router name.
all	Specifies all virtual routers.

Default

Enabled on all virtual routers.

Usage Guidelines

Use this command to disable SNMP access on any or all virtual routers.

When SNMP access is disabled on a virtual router, the incoming SNMP request is dropped and an EMS message is logged.

To enable SNMP access on virtual routers use the `enable snmp access vr` command.

To display the SNMP configuration and statistics on a specified virtual router, use the `show snmp vr_name` command.

Example

The following command disables SNMP access on the virtual router vr-finance:

```
disable snmp access vr vr-finance
```

History

This command was first available in ExtremeXOS 12.4.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable snmp access

```
disable snmp access {snmp-v1v2c | snmpv3}
```

Description

Selectively disables *SNMP* on the switch.

Syntax Description

snmp-v1v2c	Specifies SNMPv1/v2c access only.
snmpv3	Specifies SNMPv3 access only.

Default

Disabled.

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.

This command allows you to disable either all SNMP access, v1/v2c access only, or v3 access only.

To allow access, use the following command:

```
enable snmp access {snmp-v1v2c | snmpv3}
```

Example

The following command disables all SNMP access on the switch:

```
disable snmp access
```

History

This command was first available in ExtremeXOS 10.1.

SNMPv3 was added to ExtremeXOS 12.2. It was also included in ExtremeXOS 11.6.4 and 12.1.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable snmp community

```
disable snmp community [encrypted enc_community_name | community_name |  
                        alphanumeric-community-string | hex hex_community_name]
```

Description

Disables *SNMP* community strings on the switch.

Syntax Description

encrypted	Community name is encrypted.
<i>enc_community_name</i>	Encrypted community name.
<i>community_name</i>	Community name in ASCII format.
hex	Provide value in hexadecimal.
<i>hex_community_name</i>	Community name in hexadecimal.
<i>alphanumeric-community-string</i>	Specifies the SNMP community string name.

Default

N/A.

Usage Guidelines

This command allows the administrator to disable an snmp community. It sets the row status of the community to NotInService. When disabled, SNMP access to the switch using the designated community is not allowed.

Example

The following command disables the community string named extreme:

```
disable snmp community hex 61:01
```

History

This command was first available in ExtremeXOS 12.1.

The **hex** keyword and *hex_community_name* variable were added in ExtremeXOS 15.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable snmp notification-log

```
disable snmp notification-log [ default | name | hex hex_name | all ]
```

Description

Controls the administrative state of a log.

Syntax Description

default	The default log.
<i>name</i>	Specifies the name of the log.
hex	Provide value in hexadecimal.
<i>hex_name</i>	Name of the log in hexadecimal.
all	Specifies all logs.

Default

Disabled.

Usage Guidelines

Use this command to control the administrative state of a log.

Example

The following example disables *nmslog1*:

```
disable snmp notification-log hex 01:02
```

History

This command was first available in ExtremeXOS 15.5.

The **default** and **hex** keywords and *hex_name* variable were added in ExtremeXOS 15.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable snmp trap l3vpn

```
disable snmp trap l3vpn {vr name}
```

Description

This command disables Layer 3 VPN MIB notification traps for the child VPN VRFs of the specified VR.

Syntax Description

<i>vr-name</i>	Specifies the name of the parent VR where this RFC 4382 scalar is applied. If <i>vr-name</i> is not provided, then this command is applied to the VR in the current context.
----------------	--

Default

Disabled.

Usage Guidelines

None.

Example

The following example disables SNMP traps for Layer 3 VPNs on the default VR:

```
disable snmp traps l3vpn vr vr-default
```

History

This command was first available in ExtremeXOS 12.6.0-BGP.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable snmp traps

```
disable snmp traps
```

Description

Prevents SNMP traps from being sent from the switch.

Syntax Description

This command has no arguments or variables.

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.

To view if SNMP traps are being sent from the switch, use the `show management` command. The `show management` command displays information about the switch including the enabled/disabled state of SNMP traps being sent.

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 ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable snmp traps bfd

```
disable snmp traps bfd session down | session-up
```

Description

This command disables session up/down trap reception for BFD.

Syntax Description

snmp	Configure <i>SNMP</i> specific settings.
traps	Configure SNMP Trap generation settings.
bfd	BFD-specific traps.
session-down	Generate trap when BFD session goes down.
session-up	Generate trap when BFD session goes up.

Default

Both session-down and session-up.

Usage Guidelines

Use this command to disable trap reception for BFD session up/down.

Example

The following command will disable trap generation for BFD session down events.

```
# disable snmp traps bfd session-down
```

History

This command was first available in ExtremeXOS 15.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable snmp traps configuration

```
disable snmp traps configuration [save | change]
```

Description

Disables sending *SNMP* trap when saving or changing the switch configuration.

Syntax Description

configuration	Sends SNMP trap for switch configuration.
save	Disables SNMP trap when switch configuration is saved (default is disabled).
change	Disables SNMP trap when switch configuration is changed (default is disabled).

Default

The default is that SNMP traps are disabled for switch configuration changes/saves.

Example

The following example disables SNMP traps for switch configuration saves:

```
disable snmp traps configuration save
```

History

This command was first available in ExtremeXOS 22.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches..

disable snmp traps fdb mac-tracking

```
disable snmp traps fdb mac-tracking
```

Description

Disables SNMP trap generation when MAC-tracking events occur for a tracked MAC address.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

None.

Example

The following example disables SNMP traps for MAC-tracking events:

```
disable snmp traps fdb mac-tracking
```

History

This command was first available in ExtremeXOS 12.3.

Platform Availability

This command is available on all platforms.

disable snmp traps identity-management

```
disable snmp traps identity-management
```

Description

Disables the identity management feature to send *SNMP* traps for low memory conditions.

Syntax Description

This command has no arguments or variables.

Default

No traps are sent.

Usage Guidelines

None.

Example

The following command disables the identity management SNMP trap feature:

```
disable snmp traps identity-management
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable snmp traps l2vpn

```
disable snmp traps l2vpn
```

Description

Disables SNMP traps associated with Layer 2 VPNs for all MPLS configured VLANs.

Syntax Description

This command has no arguments or variables.

Default

All Layer 2 VPN traps are disabled.

Example

The following command disables SNMP traps associated with Layer 2 VPNs:

```
disable snmp traps l2vpn
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

disable snmp traps l3vpn

```
disable snmp traps l3vpn {vr vr_name}
```

Description

Use this command to turn off *SNMP* trap support for L3 VPN.

Syntax Description

<i>vr_name</i>	Specifies the name of the parent VR where this RFC 4382 scalar is applied. If <i>vr_name</i> is not provided, then this command is applied to the VR in the current context.
----------------	--

Default

Enabled.

Usage Guidelines

Use this command to disable L3VPN SNMP traps.

Example

The following example disables L3 VPN SNMP traps support on the switch:

```
disable snmp traps l3vpn vr vr-default
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable snmp traps lldp

```
disable snmp traps lldp {ports [all | port_list]}
```

Description

Disables the sending of *LLDP*-specific *SNMP* traps on the specified port or ports.

Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

Disabled.

Usage Guidelines

If you do not specify any ports, the system stops sending LLDP traps from all ports on the switch.

Example

The following example disables sending LLDP SNMP traps on all switch ports:

```
disable snmp traps lldp ports all
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable snmp traps lldp-med

```
disable snmp traps lldp-med {ports [all | port_list]}
```

Description

Disables the sending of LLDP MED-specific SNMP traps on the specified port or ports.

Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

Disabled.

Usage Guidelines

If you do not specify any ports, the system stops sending LLDP MED traps from all ports on the switch.

Example

The following example disables sending LLDP MED SNMP traps on all switch ports:

```
disable snmp traps lldp-med ports all
```

History

This command was first available in ExtremeXOS 11.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable snmp traps mpls

```
disable snmp traps mpls
```

Description

Disables *SNMP* traps associated with *MPLS* for all MPLS configured VLANs.

Syntax Description

This command has no arguments or variables.

Default

All MPLS traps are disabled.

Example

The following command disables SNMP traps associated with MPLS:

```
disable snmp traps mpls
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

disable snmp traps ospf

```
disable snmp traps ospf
```

Description

Disables the *OSPF* module from sending traps on various OSPF events.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

None.

Example

The following command disables the OSPF process:

```
disable snmp traps ospf
```

History

This command was first available in ExtremeXOS 12.6.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the *Switch Engine 32.3 Feature License Requirements* document.

disable snmp traps ospfv3

```
disable snmp traps ospfv3
```

Description

Disables the transmission of *OSPFv3 SNMP* notifications.

Syntax Description

ospfv3	OSPFv3-related traps.
---------------	-----------------------

Default

The default is disabled.

Example

The following example disables the transmission of OSPFv3 SNMP notifications:

```
disable snmp traps ospfv3
```

History

This command was first available in ExtremeXOS 22.1.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

disable snmp traps port-up-down ports

```
disable snmp traps port-up-down ports [port_list | all]
```

Description

Disables port up/down trap reception for specified ports.

Syntax Description

port_list	Specifies one or more ports or slots and ports.
all	Specifies all ports on the switch.

Default

Enabled.

Usage Guidelines

Use this command to stop receiving [SNMP](#) trap messages when a port transitions between being up and down.

Example

The following command stops ports 3, 5, and 12 through 15 on a stand-alone switch from receiving SNMP trap messages when the port goes up/down:

```
disable snmp traps port-up-down ports 3,5,12-15
```

History

This command was first available in ExtremeXOS 10.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable snmpv3

disable snmpv3 **default-group**

Description

Disables SNMPv3 default-group access on the switch.

Syntax Description

default-group	Specifies SNMPv3 default-group.
----------------------	---------------------------------

Default

Enabled.

Usage Guidelines

This command is used to disable SNMPv3 default-group access.

Disabling SNMPv3 default-group access removes access to default-users and user-created users who are part of the default-group. The user-created authenticated SNMPv3 users (who are part of a user-created group) are able to access the switch.

The default groups are: admin, initial, v1v2c_ro, v1v2c_rw.

Example

The following command disables the default group on the switch:

```
disable snmp default-group
```

History

This command was available in ExtremeXOS 12.2.

It was also included in ExtremeXOS 11.6.4 and ExtremeXOS 12.1.2.

The default-user option was removed in ExtremeXOS 22.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable snmpv3 community

```
disable snmpv3 community [ community_index | hex hex_community_index ]
```

Description

This command disables a community entry specified by the community index.

Syntax Description

community_index	Community index in ASCII.
hex	Provide value in hexadecimal.
hex_community_index	Community index in hexadecimal.

Default

Enabled.

Usage Guidelines

This command is used to disable a community entry specified by the community index.

Example

```
disable snmpv3 community hex 61:62:63:64
```

History

This command was available in ExtremeXOS 15.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

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 ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable ssh2

```
disable ssh2
```

Description

Disables the SSH2 server for incoming SSH2 sessions to switch.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

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

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

Example

The following command disables the SSH2 server:

```
disable ssh2
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable stacking

```
disable stacking {node-address node-address}
```

Description

This command disables the stacking on one or all nodes in the stack topology.

Syntax Description

node-address	Specifies the MAC address of a node in the stack. To view the MAC addresses for all nodes in a stack, enter the <code>show stacking</code> command.
---------------------	---

Default

Default value is stacking disabled.

Usage Guidelines

If you do not specify the node-address, stacking is disabled on all nodes in the stack topology.

If the node-address parameter is present, stacking is disabled on the node with the specified node-address. This is the MAC address assigned to the stackable by the factory.

A node in the stack topology that is disabled for stacking does not forward the customer's data through its stacking links and does not become a member of the active topology.

A disabled node becomes its own master and processes and executes its own configuration independently.

When this command is executed successfully, the following message appears:

This command will take effect at the next reboot of the specified node(s) .

Use `show stacking configuration` command to see the current configuration of the stack. Verify the flags in `show stacking configuration` output to confirm that stacking is disabled on the specified node(s).

Example

The following example disables stacking on an 8 node stack:

```
* Switch.3 # disable stacking
This command will take effect at the next reboot of the specified node(s) .
```

The following example disables stacking on the node with the factory assigned MAC address 00:04:96:26:6b:ed:

```
* Switch.3 # disable stacking node-address 00:04:96:26:6b:ed
This command will take effect at the next reboot of the specified node(s) .
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available with all licenses and platforms that support the SummitStack feature. For information about which licenses and platforms support the SummitStack feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

disable stacking-support

disable stacking-support

Description

This command disables the stacking-support option on a switch with dual-purpose hardware.

Syntax Description

This command does not have additional syntax.

Default

Disabled.

Usage Guidelines

The Stacking-Support Option Control column in [Table 18](#) on page 1357 displays Yes in the rows for switch configurations for which you can disable the stacking-support option.

After you disable the stacking-support option, you must reboot the switch to activate the configuration change.

If you disable the stacking-support option on a switch and reboot, stacking communication stops and the data ports listed in [Table 18](#) on page 1357 use Ethernet protocols instead of stacking protocols.

Example

To disable the stacking ports, enter the following command:

```
# disable stacking-support
This setting will take effect at the next reboot of this switch.
```

History

This command was first available in ExtremeXOS 12.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable stpd

disable stpd {*stpd_name*}

Description

Disables the STP protocol on a particular STPD or for all STPDs.

Syntax Description

<i>stpd_name</i>	Specifies an STPD name on the switch.
------------------	---------------------------------------

Default

Disabled.

Usage Guidelines

After you have created the STPD with a unique name, the keyword stpd is optional.

If you want to disable the STP protocol for all STPDs, do not specify an STPD name.

In an *MSTP* environment, you cannot delete or disable a CIST if any of the MSTIs are active in the system.

Example

The following command disables an STPD named `purple_st`:

```
disable stpd purple_st
```

The following command disables the STP protocol for all STPDs on the switch:

```
disable stpd
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable stpd auto-bind

```
disable stpd stpd_name auto-bind [ {vlan} vlan_name | vlan vlan_list]
```

Description

Disables the ability to automatically add ports to an *STPD* when they are added to a member *VLAN*.

Syntax Description

<i>stpd_name</i>	Specifies an STPD name on the switch.
<i>vlan_name</i>	Specifies the name of a member VLAN with autobind enabled.
<i>vlan_list</i>	Specifies a VLAN list of IDs.

Default

The autobind feature is disabled on user-created STPDs. The autobind feature is enabled on the default VLAN that participates in the default STPD S0.

Usage Guidelines



Note
Ports already in the STPD remain in that domain (as if they were added manually).

If you create an STPD and a VLAN with unique names, the keywords **stpd** and **vlan** are optional.

Ports added to the STPD automatically when autobind is enabled are not removed when autobind is disabled. The ports are present after a switch reboot.

To view *STP* configuration status of the ports in a VLAN, use the following command:

```
show {vlan} {vlan_name | vlan_list} stpd
```

Example

The following example disables autobind on an STPD named s8:

```
disable stpd s8 auto-bind v5
```

History

This command was first available in ExtremeXOS 10.1.

The *vlan_list* variable was added in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable stpd ports

Disables *STP* on one or more ports for a given *STPD*.

```
disable stpd stpd_name ports [all | port_list]
```

Syntax Description

<i>stpd_name</i>	Specifies an STPD name on the switch.
all	Specifies all ports for a given STPD.
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

Enabled.

Usage Guidelines

If you create the STPD with a unique name, the keyword stpd is optional.

Disabling STP on one or more ports puts those ports in the forwarding state; all BPDUs received on those ports are disregarded and dropped.

Use the all keyword to specify that all ports of a given STPD are disabled.

Use the `port_list` parameter to specify a list of ports of a given STPD are disabled.

If you do not use the default STPD, you must create one or more STPDs and configure and enable the STPD before you can use the `disable stpd ports` command.

Example

The following command disables slot 2, port 4 on an STPD named Backbone_st:

```
disable stpd backbone_st ports 2:4
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable stpd rapid-root-failover

```
disable stpd stpd_name rapid-root-failover
```

Description

Disables rapid root failover for *STP* recovery times.

Syntax Description

<i>stpd_name</i>	Specifies an <i>STPD</i> name on the switch.
------------------	--

Default

Disabled.

Usage Guidelines

This command is applicable for STPDs operating in 802.1D.

After you have created the STPD with a unique name, the keyword `stpd` is optional.

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 ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable switch bluetooth

```
disable switch bluetooth {discovery | pairing }
```

Description

Disables Bluetooth capability on a switch.

Syntax Description

switch	Designates disabling switch capabilities.
bluetooth	Designates disabling Bluetooth capabilities on a switch.
discovery	Disables discoverable mode of the switch. Default is enabled.
pairing	Disables pairing ability with other Bluetooth-capable devices. Default is enabled.

Default

By default, discovery and pairing modes are enabled.

Usage Guidelines

Using the command with no options disables Bluetooth capability on the switch. The **discovery** and **pairing** options disable discoverable mode and pairing ability, respectively.

To enable Bluetooth capabilities, use the `enable switch bluetooth {discovery | pairing }` command.

To view Bluetooth and discovery/pairing status, use the `show switch bluetooth [statistics | inventory]` command.

Example

The following example disables Bluetooth capability on a switch:

```
# disable switch bluetooth
```

The following example disables discovery mode on a switch:

```
# disable switch bluetooth discovery
```

History

This command was first available in ExtremeXOS 30.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable switch locally-administered-address

```
disable switchlocally-administered-address
```

Description

Disables the switch from generating locally administered per-port MAC addresses.

Syntax Description

This command has no arguments or variables.

Default

This feature is disabled by default.

Usage Guidelines

ExtremeXOS switches do not use a unique per-port MAC address when transmitting bridge protocol data units (BPDUs). As a result, switch management can become inaccessible when switch MAC addresses are learned on the wrong L2 path (corresponding to a blocking port). This command allows you to disable the switch from generating locally administered MAC addresses.

Example

The following example disables the switch from generating locally administered MAC addresses:

```
disable switch locally-administered-address
```

History

This command was first available in ExtremeXOS 22.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable switch usb

`disable switch usb`

Description

Disables use of the switch's USB port.

Syntax Description

usb	Specifies USB port on switch.
------------	-------------------------------

Default

Enabled by default.

Usage Guidelines

This command requires a reboot to take effect.

Stack support is not available. You need to run this command individually on each node in a stack.

Running `unconfigure switch all` removes this USB setting and returns to the default of enabled.

Example

The following example disables use of the USB port:

```
disable switch usb
This setting will take effect at the next system reboot.
```

History

This command was first available in ExtremeXOS 22.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable syslog

`disable syslog`

Description

Disables logging to all remote syslog server targets.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

Disables logging to all remote syslog server targets, not to the switch targets. This setting is saved in FLASH, and will be in effect upon boot up.

Example

The following command disables logging to all remote syslog server targets:

```
disable syslog
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable subvlan-proxy-arp vlan

```
disable subvlan-proxy-arp vlan [vlan-name | all]
```

Description

Disables the automatic entry of subVLAN information in the proxy ARP table.

Syntax Description

<i>vlan-name</i>	Specifies a superVLAN name.
all	Specifies all VLANs.

Default

Enabled.

Usage Guidelines

To facilitate communication between subVLANs, by default, an entry is made in the IP ARP table of the superVLAN that performs a proxy ARP function. This allows clients on one subVLAN to communicate with clients on another subVLAN. In certain circumstances, intra-subVLAN communication may not be desired for isolation reasons.



Note

The isolation option works for normal, dynamic, ARP-based client communication.

Example

The following example disables the automatic entry of subVLAN information in the proxy ARP table of the superVLAN "vsuper":

```
disable subvlan-proxy-arp vlan vsuper
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable tacacs

```
disable tacacs
```

Description

Disables TACACS+ authentication.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command disables TACACS+ authentication for the switch:

```
disable tacacs
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

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 ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable tacacs-authorization

```
disable tacacs-authorization
```

Description

Disables TACACS+ authorization.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This disables 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 ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable tech-support collector

```
disable tech-support collector
```

Description

Disables the tech support feature.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command disables the tech-support feature. In the ExtremeXOS 15.4 release, the feature is disabled by default. When the feature is disabled, the previous scheduled reports are canceled, and the bootup event and critical severity events are ignored.

Example

The following command disables the tech-support feature:

```
disable tech-support collector
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable telnet

```
disable telnet
```

Description

Disables external 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.



Note

Telnet sessions between the nodes of a stack are not affected by this command.

Example

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

```
disable telnet
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable tunnel

```
disable {tunnel} tunnel_name
```

Description

Allows GRE tunnels to be disabled.

Syntax Description

<i>tunnel_name</i>	GRE tunnel name.
--------------------	------------------

Default

Enabled.

Usage Guidelines

Use this command to disable GRE tunnels.

Example

This example disables the tunnel named "myGREtunnel":

```
disable myGREtunnel
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable twamp reflector

```
disable twamp reflector {restrict}
```

Description

This command disables the Session-Reflector.

Syntax Description

restrict	Restricts only TWAMP control sessions to create test sessions and reflector does not respond to TWAMP-test packets that do not match a test session created by a control session.
-----------------	---

Default

N/A.

Usage Guidelines

If you disable the Session-Reflector, the application terminates all current TWAMP test sessions. If you specify the **restrict** keyword, only TWAMP control sessions may create test sessions and the reflector will not respond to TWAMP-test packets that do not match a test session created by a control session.

History

This command was first available in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable twamp server

```
disable twamp server
```

Description

This command disables the TWAMP server.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

If you disable the TWAMP server, all current TWAMP control sessions are terminated and any test sessions set up by the control sessions are deleted.

History

This command was first available in ExtremeXOS 16.1.

Platform Availability

The command is available on all platforms.

disable udp-echo-server

```
disable udp-echo-server {vr vrid}
```

Description

Disables UDP echo server support.

Syntax Description

<i>vrid</i>	Specifies a VR or VRF.
-------------	------------------------

Default

Disabled.

Usage Guidelines

UDP echo packets are used to measure the transit time for data between the transmitting and receiving end.

Example

The following example disables UDP echo server support:

```
disable udp-echo-server
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable upm profile

```
disable upm profile profile-name
```

Description

Disables the use of the specified Universal Port profile on the switch.

Syntax Description

<i>profile-name</i>	Specifies the UPM profile to be disabled.
---------------------	---

Default

A UPM profile is enabled by default.

Example

The following command disables a UPM profile called sample_1 on the switch:

```
disable upm profile sample_1
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the Universal Port feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

disable virtual-network remote-endpoint vxlan

```
disable virtual-network remote-endpoint vxlan [ ipaddress ipaddress {vr  
  vr_name} | all ]
```

Description

Disables a VXLAN remote endpoint.

Syntax Description

virtual-network	Virtual overlay network.
remote-endpoint	Remote tunnel endpoint information.
vxlan	VXLAN virtual networks remote endpoint.
ipaddress	Specifies an IP address of a remote endpoint.
<i>ipaddress</i>	Specifies the IP address of the desired remote endpoint.
vr	Specifies a VR/VRF instance the remote endpoint is associated with.
<i>vr_name</i>	Specifies the desired existing VR/VRF instance the remote endpoint is associated with. Default is VR-Default.
all	Specifies all remote tunnel endpoints.

Default

If a VR is not specified, VR-Default is the VR.

Usage Guidelines

Extreme Loop Recognition Protocol (ELRP) detects loops across VXLAN tunnels. If a loop is detected across the tunnel, ELRP takes down the VXLAN remote endpoint. You can use this command to disable a remote endpoint manually.

Example

The following example disables the remote endpoint at 100.1.1.1 on VR-Default (not specified, command default):

```
# disable virtual-network remote-endpoint vxlan ipaddress 100.1.1.1
```

History

This command was first available in ExtremeXOS 22.4.

Platform Availability

This command is available on ExtremeSwitching 5420 and 5520 series switches.

disable virtual-router

```
disable virtual-router vrf-name
```

Description

Disables a VRF.

**Note**

This command is only applicable for VRFs.

Syntax Description

<i>vrf-name</i>	Specifies the name of the VRF.
-----------------	--------------------------------

Default

Enabled.

Usage Guidelines

When you disable a VRF, the software does the following:

- Disables Layer 3 protocols.
- Marks static routes as inactive and removes them from the hardware forwarding tables.
- Flushes the IP ARP and IPv6 neighbor-discovery caches.

Example

The following example disables VRF "vrf1":

```
disable virtual-router vrf1
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable vlan

```
disable [ {vlan} vlan_name | vlan vlan_list]
```

Description

Use this command to disable the specified VLAN.

Syntax Description

<i>vlan_name</i>	Specifies the VLAN you want to disable.
<i>vlan_list</i>	Specifies the VLAN list of IDs to disable.

Default

Enabled.

Usage Guidelines

This command allows you to administratively disable specified VLANs. The following guidelines apply to working with disabling VLANs:

- Disabling a VLAN stops all traffic on all ports associated with the specified VLAN.
- You cannot disable a VLAN that is running Layer 2 protocol control traffic for protocols such as *EAPS*, *STP*, or *ESRP*.

When you attempt to disable a VLAN running Layer 2 protocol control traffic, the system returns a message similar to the following:

```
VLAN accounting cannot be disabled because it is actively used by an L2 Protocol
```

- You can disable the default VLAN; ensure that this is necessary prior to disabling the default VLAN.
- You cannot disable the management VLAN.
- You cannot bind Layer 2 protocols to a disabled VLAN.
- You can add ports to or delete ports from a disabled VLAN.



Caution

Disabling the Mgmt VLAN disables access to the Ethernet Management port on a switch (`disable vlan Mgmt`).

Example

The following example disables the VLAN named "accounting":

```
disable vlan accounting
```

History

This command was first available in ExtremeXOS 11.4.

The ability to add ports to a disabled VLAN was added in ExtremeXOS 12.5.

The `vlan_list` option was added in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable vm autostart

```
disable vm vm_name autostart
```

Description

Disables automatic start-up of guest virtual machines (VMs).

Syntax Description

vm	Designates a virtual machine.
<i>vm_name</i>	Specifies the VM name.
autostart	Specifies disabling automatic start-up of the specified VM. Default is disabled.

Default

By default, automatic start-up is disabled.

Usage Guidelines

This command disables automatically starting up a specific VM when the system starts.

You must reboot the switch for this command to take effect.

To enable automatic start-up, use the command `enable vm vm_name autostart`.

The Integrated Application Hosting (IAH) feature requires the Solid State Storage Device SSD-120.

Example

The following example disables automatic start-up of VM "vm1":

```
# disable vm vm1 autostart
```

History

This command was first available in ExtremeXOS 30.3.

Platform Availability

This command is available on all platforms that support the Integrated Application Hosting (IAH) feature and have a Core or Premier license installed. For a list of platforms that support the IAH feature and for information about licenses, see the [Switch Engine 32.3 Feature License Requirements](#).

disable vm-tracking dynamic-vlan ports

```
disable vm-tracking dynamic-vlan ports port_list
```

Description

This command disables VM-tracking dynamic VLAN on specific ports.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

Use this command to disable VM-tracking dynamic VLAN on specific ports. The ALL option is not supported because VM-tracking dynamic VLAN should not be enabled on a switch's uplink port.

Example

This example disables VM-tracking dynamic VLAN on port 2:1:

```
# disable vm-tracking dynamic-vlan ports 2:1
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable vm-tracking

```
disable vm-tracking
```

Description

Disables the *Extreme Network Virtualization (XNV)* feature on the switch.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command disables the XNV feature, which tracks virtual machines (VMs) that connect to the switch.



Note

When the VM tracking feature is disabled, file synchronization with the FTP server stops.

Example

The following command disables the XNV feature:

```
# disable vm-tracking
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable vm-tracking ports

```
disable vm-tracking ports port_list
```

Description

Disables the XNV feature on the specified ports.

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
------------------	---

Default

Disabled.

Usage Guidelines

This command disables VM tracking on the specified ports.

Example

The following command disables VM tracking on port 2:1:

```
# disable vm-tracking ports 2:1
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable vman cep egress filtering ports

```
disable vman cep egress filtering ports {port_list | all}
```

Description

Disables the egress filtering of CVIDs that are not configured in the CVID map for a CEP.

Syntax Description

<i>port_list</i>	Specifies a list of ports.
all	Specifies all switch ports.

Default

Egress CVID filtering is disabled.

Usage Guidelines

To view the configuration setting for the egress CVID filtering feature, use the `show ports information` command.



Note

When CVID egress filtering is enabled, it reduces the maximum number of CVIDs supported on a port. The control of CVID egress filtering applies to fast-path forwarding. When frames are forwarded through software, CVID egress filtering is always enabled.

Example

The following example disables egress CVID filtering on port 1:

```
disable vman cep egress filtering port 1
```

History

This command was first available in ExtremeXOS 12.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable vpex

```
disable vpex
```

Description

Disables VPEX mode for using bridge port extenders (BPEs).

Syntax Description

vpex	Specifies Virtual Port Extender (VPEX).
-------------	---

Default

N/A.

Usage Guidelines

Disabling VPEX mode removes all BPE slot number assignments made using `configure vpex ports port_list slot slot_num`. A reboot of the switch is required for this command to take effect.

Example

The following example disables VPEX mode:

```
# disable vpex
```

History

This command was first available in ExtremeXOS 22.5.

Platform Availability

This command is available on ExtremeSwitching 5420 and 5520 series switches.

disable vpex auto-configuration

```
disable vpex auto-configuration
```

Description

Disables automatic configuration of the Extended Edge Switching architecture (controlling bridge (CB) and bridge port extenders (BPEs)).

Syntax Description

vpex	Specifies Virtual Port Extender (VPEX).
auto-configuration	Specifies disabling automatic configuration of the Extended Edge Switching architecture.

Default

Disabled.

Usage Guidelines

Auto-configuration allows the controlling bridge switch to detect new BPEs connected to ports not configured as cascade ports, and automatically configure cascade ports, LAG membership, ports, and extended slots. This command disables this auto-configuring capability.

To disable auto-configuration, you must first enter VPEX mode (see [enable vpex](#) on page 2353).

Example

The following example disables auto-configuration mode:

```
# disable vpex auto-configuration
```

History

This command was first available in ExtremeXOS 22.6.

Platform Availability

This command is available on ExtremeSwitching 5420 and 5520 series switches.

disable vpex auto-upgrade

disable **vpex auto-upgrade**

Description

Disables automatic upgrading on Extended Edge Switching topologies.

Syntax Description

vpex	Specifies Virtual Port Extender (VPEX).
auto-upgrade	Specifies that the controlling bridge (CB) automatically upgrades bridge port extender (BPE) slots in mode (default is enabled).

Default

Automatic upgrading is enabled by default.

Usage Guidelines

Automatic upgrading can occur only when both CBs in the MLAG have the same BPE xmod versions installed, and only after all slots are synchronized between the CBs.

To disable automatic upgrading, you must first enter VPEX mode (see [enable vpex](#) on page 2353). To view the status of automatic upgrading, use the command `show vpex`.

Example

The following example disables automatic upgrading:

```
# disable vpex auto-upgrade
```

History

This command was first available in ExtremeXOS 30.5.

Platform Availability

This command is available on ExtremeSwitching 5420 and 5520 series switches.

disable vpls

```
disable vpls [vpls_name | all]
```



Note

This command has been replaced with the following command: `disable l2vpn [vpls [vpls_name | all] | vpws [vpws_name | all]]`.

This command is still supported for backward compatibility, but it will be removed from a future release, so we recommend that you start using the new command.

Description

Disables the VPLS instance specified by `vpls_name`.

Syntax Description

<code>vpls_name</code>	Identifies the VPLS within the switch (character string)
all	Specifies all VPLS.

Default

All newly created VPLS instances are enabled.

Usage Guidelines

This command disables the VPLS instance specified by `vpls_name`. When a VPLS instance is disabled, all sessions to its configured peers are terminated. Any locally attached service VLAN/VMAN is immediately isolated from other devices residing in the VPN. If this is an H-VPLS core node, then all spoke nodes connected to this peer are isolated unless redundant core access is configured.

Example

The following example disables the VPLS named "myvpls":

```
disable vpls myvpls
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

disable vpls fdb mac-withdrawal

```
disable vpls fdb mac-withdrawal
```



Note

This command has been replaced with the following command: `disable l2vpn vpls fdb mac-withdrawal` .

This command is still supported for backward compatibility, but it will be removed from a future release, so we recommend that you start using the new command.

Description

Disables the VPLS MAC address withdrawal capability.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

When disabled, the switch does not send MAC address withdrawal messages. If a MAC address withdrawal message is received from another VPLS peer, the local VPLS peer processes the message and withdraws the specified MAC addresses from its [FDB](#), regardless of the MAC address withdrawal configuration.

Example

The following command disables MAC address withdrawal:

```
disable vpls fdb mac-withdrawal
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available only on the platforms that support *MPLS* as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

disable vpls health-check vccv

```
disable vpls [vpls_name | all] health-check vccv
```



Note

This command has been replaced with the following command: `disable l2vpn [vpls [vpls_name | all] | vpws [vpws_name | all]] health-check vccv`.

This command is still supported for backward compatibility, but it will be removed from a future release, so we recommend that you start using the new command.

Description

Disables the VCCV health check feature on one or all VPLS instances on the local node.

Syntax Description

<code>vpls_name</code>	Identifies the VPLS for which health check is to be disabled.
all	Specifies that health check is to be disabled on all VPLS instances on the local node.

Default

Health check is disabled.

Usage Guidelines

None.

Example

The following command disables the health check feature on the VPLS instance myvpls:

```
disable vpls myvpls health-check vccv
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available only on the platforms that support *MPLS* as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

disable vpls service

```
disable vpls [vpls_name | all] service
```



Note

This command has been replaced with the following command: `disable l2vpn [vpls [vpls_name | all] | vpws [vpws_name | all]] service`.

This command is still supported for backward compatibility, but it will be removed from a future release, so we recommend that you start using the new command.

Description

Disables the configured VPLS services for the specified *vpls_name*.

Syntax Description

<i>vpls_name</i>	Identifies the VPLS within the switch (character string).
all	Specifies all VPLS.

Default

Enabled.

Usage Guidelines

When services are disabled, the VPLS is removed from all peer sessions. The keyword **all** disables services for all VPLS instances.

Example

The following command disables the configured VPLS services for the specified VPLS:

```
disable vpls myvpls service
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support *MPLS* as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

disable vrrp group

```
disable vrrp group group_name {configuration | members}
```

Description

This command disables group mode on member VRs so that they can operate in individual VR mode.

Syntax Description

group	Form a group of <i>VRRP</i> VRs to operate in high-scale mode.
<i>group_name</i>	Specifies the VRRP group name.
configuration	Removes group configuration on individual VRs (default).
members	Disables all VRs that are members of the group.

Default

If you do not specify, group configuration is removed from individual VRs.

Example

The following example disables administratively all member VRs of the group. This may be useful for debugging issues:

```
disable vrrp group ExtremeNet members
```

History

This command was first available in ExtremeXOS 22.2.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the VRRP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

disable vrrp vrid

```
disable vrrp {vlan [vlan_name | vlan_list] vrid [vridval | vrid_list]}
```

Description

Disables a specific VRRP instance or all VRRP instances.

Syntax Description

<i>vlan_name</i>	Specifies the name of a VRRP <u>VLAN</u> .
<i>vlan_list</i>	VLAN list (1-4,094).
<i>vridval</i>	Specifies the VRID for the VRRP instance. To display the configured VRRP router instances, enter the show vrrp command.
<i>vrid_list</i>	List of virtual router IDs (1-255).

Default

N/A.

Usage Guidelines

This disables a specific VRRP instance on the switch. If no VRRP VLAN is specified, all VRRP instances on the switch are disabled.

Example

The following command disables all VRRP instances on the switch:

```
disable vrrp
```

History

This command was first available in ExtremeXOS 10.1.

VLAN and VR list options added in ExtremeXOS 22.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the VRRP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

disable watchdog

```
disable watchdog
```

Description

Disables the system watchdog timer.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

The watchdog timer monitors the health of the switch hardware and software events. For example, the watchdog timer reboots the switch if the system cannot reset the watchdog timer. This can be caused by a long CPU processing loop, any unhandled exception, or a hardware problem with the communication channel to the watchdog. In most cases, if the watchdog timer expires, the switch captures the current CPU status and posts it to the console and the system log. In some cases, if the problem is so severe that the switch is unable to perform any action, the switch reboots without logging any system status information prior to reboot.

This command takes affect immediately.

The watchdog settings are saved in the configuration file.

To display the watchdog state of your system, use the `show switch` command.

Example

The following command disables the watchdog timer:

```
disable watchdog
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable web http

```
disable web http
```

Description

Disables the hypertext transfer protocol (HTTP) access to the switch on the default port (80).

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

Use this command to disallow users from connecting with HTTP. Disabling HTTP access forces a user to use a secured HTTPS connection if web HTTPS is enabled.

Use the following command to enable web HTTPS:

```
enable web https
```

Example

The following command disables HTTP on the default port:

```
disable web http
```

History

This command was first available in the ExtremeXOS 11.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable web https

```
disable web https
```

Description

Disables the secure socket layer (SSL) access to the switch on the default port (443).

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

Use this command to disable SSL before changing the certificate or private key.

Example

The following command disables SSL on the default port:

```
disable web https
```

History

This command was first available in the ExtremeXOS 11.2 and supported with the SSH module.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable cli xml-mode

```
disable cli xml-mode
```

Description

Disables XML configuration mode on the switch.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

Use this command to disable the XML configuration mode on the switch. XML configuration mode is not supported for end users.

See the command:


```
enable xml-mode
```

Example

The following command disables XML configuration mode on the switch:

```
disable cli xml-mode
```

History

This command was first available in an ExtremeXOS 11.2.

The **cli** keyword was added for syntax consistency in ExtremeXOS 30.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

disable msrp ports

```
disable msrp ports [port_list | all]
```

Description

Disables MSRP on the ports listed in the command after the keyword **ports**.

Syntax Description

msrp	Multiple Stream Registration Protocol.
<i>port_list</i>	Port list separated by a comma or "-".
all	All ports.

Default

Disabled.

Usage Guidelines

Use this command to disable MSRP in the ports listed or all ports.

Example

```
disable msrp ports all
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on all platforms if the AVB feature pack license is installed on the switch.

download bootrom

```
download bootrom [[ipaddress | hostname] filename {{vr} vrname} {block-size block_size}] {slot slotid} {install} {reboot}}
```

Description

Downloads a BootROM image after the switch has booted.

The downloaded image replaces the BootROM in the onboard FLASH memory.

Syntax Description

<i>ipaddress</i>	Specifies the IP address of the TFTP server.
<i>hostname</i>	Specifies the host name of the TFTP server. Use of the <i>hostname</i> option requires that DNS be enabled.
<i>filename</i>	The name of the bootROM file (.xtr extension).
<i>vrname</i>	Specifies the name of the virtual router. Note: User-created VRs are supported only on the platforms listed for this feature in the Switch Engine 32.3 Feature License Requirements document.
<i>block_size</i>	Specifies the data block size, excluding TFTP header. Data block size ranges from 24-65,000 bytes.
<i>slotid</i>	This parameter is available only on the SummitStacks. On a SummitStack, the <i>slotid</i> specifies the slot number of the node on which the image should be downloaded.
install	Specifies installing the system image after download.
reboot	Specifies rebooting after installation.

Default

The default block size is 1,400 bytes.

Usage Guidelines

Upgrade the BootROM image only when asked to do so by an Extreme Networks technical representative.

The BootROM image file is designated with a .xtr file extension.

Prior to downloading the BootROM image on the switch, you must download the image you received from Extreme Networks to a TFTP server on your network. You can also download the image to a USB 2.0 storage device.

When you download a BootROM image, you are prompted to install the image immediately after the download is finished. You can also use the **install** option to choose in advance to install the bootROM image. If you choose to install the image at a later time, use the `install bootrom [from-image | fname | local-file] {slot slot-number} {reboot}` command to install the software on the switch.

If this command does not complete successfully it could prevent the switch from booting. If the switch does not boot properly, some boot option functions can be accessed through a special Bootloader menu.

Displaying the BootROM Versions

To display the BootROM version for the switch, use the `show version` command.

Host Name and Remote IP Address Character Restrictions

When specifying a host name or remote IP address, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-) Permitted only for host names
- Underscore (_) Permitted only for host names
- Colon (:)

Local and Remote File Name Character Restrictions

When specifying a local or remote file name, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-)
- Underscore (_)
- Slash (/) Permitted only for remote files

SummitStack Only

You can run this command only from the master node. The file to be downloaded has to be compatible with the type of switch in the specified slot.

If you do not specify a slot number and you elect to install the image after downloading, an attempt is made to install the BootROM image on all active nodes. The BootROM image is not installed on any node if the BootROM image specified is not compatible with all active nodes.

Example

The following example downloads a BootROM image from the TFTP server "tftphost" with the file name "bootimage":

```
# download bootrom tftphost bootimage
```

History

This command was first available in ExtremeXOS 11.0.

The **slot** parameter was added to support SummitStack in ExtremeXOS 12.0.

Support for USB 2.0 storage devices was added in ExtremeXOS 12.5.3.

Block size support was added in ExtremeXOS 15.7.1.

The **memorycard** option was removed in ExtremeXOS 30.7.

The **install** and **reboot** options were added in ExtremeXOS 31.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

download image

Using TFTP: `download [url url {vr vrname} | image [active | inactive] [[hostname | ipaddress] filename {{vr} vrname} {block-size block_size}] {partition} {install {reboot}}`

To download an image to a stack: `download image [[hostname | ipaddress] filename {{vr} vrname} {block-size block_size}] {partition} {install {reboot}}`

Description

Downloads a new version of the ExtremeXOS software image or a new Fabric Engine image when changing the switch's network operating system.

The image file can be downloaded using TFTP (which is not a secure method), or SFTP and SCP2 (which are secure methods). The procedure using TFTP begins above and using SFTP/SCP2.

Syntax Description

url	Uniform Resource Locator (URL) of the file to download, which is of a supported type (for example, <code>cfg</code> , <code>lic</code> , <code>lst</code> , <code>pol</code> , <code>py</code> , Fabric Engine, <code>xmod</code> , <code>xos</code> , <code>xsf</code> , <code>xtr</code>).
<i>url</i>	Specifies the URL of the supported file (for example, <code>http://ipaddress/path.xos</code> or <code>ftp://ipaddress:port/path.xmod</code> or <code>ftp://ipaddress/some_list.lst</code>)

active	Specifies automatic determination for active (primary) partition. Note: Not applicable Fabric Engine specifying a Fabric Engine image.
inactive	Specifies automatic determination for inactive (secondary) partition. Note: Not applicable when specifying a Fabric Engine image.
<i>hostname</i>	Specifies the hostname of the TFTP server from which the image should be obtained.
<i>ipaddress</i>	Specifies the IP address of TFTP server from which the image should be obtained.
<i>filename</i>	Specifies the file name of the new image. You can use this command to change the operating system to Fabric Engine by downloading and installing a Fabric Engine image.
<i>vrname</i>	Specifies the name of the virtual router. Note: User-created VRs are supported only on the platforms listed for this feature in the Switch Engine 32.3 Feature License Requirements document.
<i>block_size</i>	Specifies the maximum block size, not including the TFTP header. The range is 24–65000 bytes.
<i>partition</i>	Specifies which partition the image should be saved to: primary or secondary. Select primary to save the image to the primary partition and secondary to save the image to the secondary partition. Note: Not applicable when specifying a Fabric Engine image.
install	Specifies installing the image after download.
reboot	Specifies rebooting after installation.

Default

Stores the downloaded image in the alternate (inactive) partition.

SFTP and SCP2 provide secure methods of downloading the ExtremeXOS software image files, *.xos or *.xmod. You can use one of three procedures:

- From the switch, running the command SCP2. connect to and “get” from a remote server. This is similar to the `download image` command.
- From outside the switch, connect to the switch that is acting as the server and “put” from the remote server. There is no TFTP equivalent for this method.
 - Using SFTP
 - Using SCP2

If you do not specify block size, the default value is 1,400 bytes.

Usage Guidelines

Prior to downloading an image on the switch, you must download the image you received from Extreme Networks to a TFTP server on your network. If your switch has a removable storage device, you can also download the image to that device.



Note

The `download image` command causes the switch to use the newly downloaded software image during the next switch reboot. To modify or reset the software image used during a switch reboot, use the `use image` command. Use this command after downloading and installing the image for it to be effective.

Specify the IP address or host name parameters to download an image from a TFTP server on the network. Use of the host name parameter requires that DNS be enabled.

When you download and install a new version of an ExtremeXOS image, the system automatically compares the currently installed bootROM image against the bootROM image contained in the new ExtremeXOS image. If the installed version is older, the system automatically upgrades to the bootROM version contained in the new ExtremeXOS image.

Core Software Images

A switch can store up to two core images: an active and inactive. When downloading a new image, you must select on which partition to install the new image. You must install the software image to the inactive partition, and must specify that partition while downloading the image to the switch.

Image Filenames

The software image file can be an `.xos` file, which contains a Switch Engine core image, or an `.xmod` file, which contains a Switch Engine modular software package.

As of ExtremeXOS 16.1, the `download` command now accepts a URL as the name of the file to download. URL protocols can be `tftp`, `http`, `ftp`. The format of a URL is:

- `http://10.10.10.1/filename.xos`
- `tftp://10.10.10.1/filename.xos`
- `ftp://10.10.10.1/filename.xmod`

In addition to accepting a URL that ends in `.xos` or `.xmod`, the URL file name can end in `.1st`. A `.1st` file contains file names at the same location as the `.1st` file URL and is downloaded/installed one after the other. The `.1st` file method defines bundles of downloads for:

- aspen, summit480 -image file size issues
- SSH installs with ExtremeXOS
- Customer files ending in `'.cfg'`, `'.xsf'`, `'.pol'`, `'.lic'`, `'.py'`, `'.ssh'`
- Other bundling that makes it easier to download with a single command

For additional installation requirements, see the sections *Installing a Core Image* and *Installing a Modular Software Package* in the [Switch Engine 32.3 User Guide](#).

Changing the Switch Network Operating System to Fabric Engine

For ExtremeSwitching Universal platforms, you can change the switch's network operating system to Fabric Engine by specifying a Fabric Engine image in *filename*.



Caution

Changing your network operating systems deletes all configuration files, debug information, logs, events, statistics, and license information of the previous network operating system.

Displaying the Software Image Versions

To display the software image version running on the switch, use the `show version` or `show switch` commands.

Host Name and Remote IP Address Character Restrictions

This section provides information about the characters supported by the switch for host names and remote IP addresses.

When specifying a host name or remote IP address, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-) Permitted only for host names
- Underscore (_) Permitted only for host names
- Colon (:)

When naming or configuring an IP address for your network server, remember the requirements listed above.

Local and Remote Filename Character Restrictions

This section provides information about the characters supported by the switch for local and remote filenames.

When specifying a local or remote filename, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-)
- Underscore (_)
- Slash (/) Permitted only for remote files

When naming a local or remote file, remember the requirements listed above.

Messages Displayed by the Switch

When you download a new image, you will see the following message:

```
Do you want to install image after downloading? (y - yes, n - no, <cr> - cancel)
```

Do one of the following:

- Enter **y** if you want to install the image after download.
- Enter **n** if you want to install the image at a later time.
- Press **[Enter]** if you want to cancel the download.

The Image Integrity Check feature was added in ExtremeXOS 16.1. The CLI output of this command is modified:

1. If the signature is verified and there is no error, there is no change to the output.
2. If the downloaded image does not have a signature, the following messages are added. This is considered as a warning, since it could be simply a downgrading. The user is given the choice to continue or quit the installation.

```
Warning: Signature Validation - Image is not digitally signed. Do you want to continue?  
(y/N)
```

If the user decides to continue, then it follows the normal installation path; if the user decides to stop here, the following message is printed and then the installation is canceled.

```
Installation canceled
```

3. If the certificate (keys) to verify the image is missing, the following messages are added. This is considered as a non-fatal and rare error, digital signature verification is bypassed. The user is given the choice to continue or quit the installation.

```
Warning: Signature Validation - Certificates missing; Image signature validation will  
be bypassed. Do you want to continue? (y/N)
```

If the user decides to continue, then it follows the normal installation path; if the user decides to stop here, the following message is printed and then the installation is canceled.

```
Installation canceled
```

4. If the certificate (keys) itself cannot be verified, the following messages are added. This is STILL considered as a non-fatal and rare error, digital signature verification is bypassed. You have the choice to continue or quit the installation.

```
Warning: Signature Validation - Certificates verification failed; Image signature  
validation will be bypassed. Do you want to continue? (y/N)
```

If you decide to continue, the normal installation continues. If you decide to stop here, the following message appears and the installation is canceled.

```
Installation canceled
```

5. If the image digital signature validation fails, the following message is added as a new reason why download fails. This is considered a fatal error like a CRC check failure, installation is terminated immediately.

```
Error: Failed to download image - Error: Image signature cannot be validated.
```

SummitStack Only

You can issue this command only from the master node.

If a slot is not specified, the image is downloaded to every node in the active topology. If a slot is specified, the image is downloaded to that slot only.

If all nodes to be downloaded are not running the same partition, the command is not executed and following message is displayed:

```
Error: all nodes do not have the same image partition selected.
```

If all nodes to be downloaded have the same partition selected but the ExtremeXOS is currently running from the selected partition, the command is not executed and the following message appears:

```
Error: the image partition selected must not be the active partition.
```

Downloading a New Image

For information about upgrading .xos and .xmod images, see the *Software Upgrade and Boot Options* section in the [Switch Engine 32.3 User Guide](#).

Example

The following example shows how the .lst file can contain filenames ending in .lst to get a list of lists (of lists etc...) from an HTTP server on 10.68.9.7 port 8080 for directory 16.1/cougar/cougar/release:

cat big.lst – big.lst contains other list file names:

- xos.lst
- xmod.lsts
- cript.lst

cat xos.lst – xos.lst contains an ExtremeXOS image:

- summitX-16.1.0.18.xos

cat xmod.lst – xmod.lst :contains a number of .xmod filenames:

- summitX-16.1.0.18-debug.xmod
- summitX-16.1.0.18-LegacyCLI.xmod
- summitX-16.1.0.18-reachnxt-1.8.1.8.xmod
- summitX-16.1.0.18-techSupport.xmod

cat script.lst – script.lst contains a number of Python scripts the user wants to download to a switch:

- jsonrpc.py
- jsontest.py
- otst.py
- ping.py
- readvr.py

A single download command downloads all of the above files.

```
# download url
http://10.68.9.7/big.lst
http://10.68.9.7/xos.lst
Downloading http://10.68.9.7/summitX-16.1.0.18.xos

Downloading to Switch.....
Installing to primary partition!

Installing to Switch.....
```

```

Image installed successfully
This image will be used only after rebooting the switch!

Image installed successfully
http://10.68.9.7:8080/xmod.lst
Downloading http://10.68.9.7/summitX-16.1.0.18-debug.xmod

Downloading to Switch.....
Installing to primary partition!

Installing to Switch.....
Image installed successfully
Downloading http://10.68.9.7/summitX-16.1.0.18-LegacyCLI.xmod

Downloading to Switch..
Installing to primary partition!

Installing to Switch.....
Legacy CLI framework was Successfully Installed !!!

Image installed successfully

Downloading http://10.68.9.7/summitX-16.1.0.18-reachnxt-1.8.1.8.xmod

Downloading to Switch...
Installing to primary partition!

Installing to Switch....
Image installed successfully
Downloading http://10.68.9.7/summitX-16.1.0.18-techSupport.xmod

Downloading to Switch..
Installing to primary partition!

Installing to Switch..
Image installed successfully
http://10.68.9.7/script.lst
http://10.68.9.7/jsonrpc.py
http://10.68.9.7/jsontest.py
http://10.68.9.7/otst.py
http://10.68.9.7/ping.py
http://10.68.9.7/readvr.py
(pacman debug) 5520-24t #

```

The following example changes the operating system to Fabric Engine:

```

# download image 10.68.9.9 voss-8.2.tgz
Do you want to install image after downloading? (y - yes, n - no, <cr> - cancel)
Downloading to Slot-1.....
.....
.....
WARNING: The specified image is for the VOSS Network Operating System and the EXOS
Network Operating System is currently running. If you continue, all configuration, logs,
and debug will be cleared and VOSS will be installed. Continue? (y - yes, n - no, <cr> -
cancel)

```

History

This command was first available in ExtremeXOS 10.1.

The **memorycard** option was added in ExtremeXOS 11.0.

Support for USB 2.0 storage devices was added in ExtremeXOS 12.5.3.

Block size support was added in ExtremeXOS 15.7.1.

The **memorycard** keyword was removed in ExtremeXOS 30.7.

Support for downloading Fabric Engine images and the **install** and **reboot** options were added in ExtremeXOS 31.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

download ssl certificate

```
download ssl ipaddress certificate {ssl-cert | trusted-ca | ocsp-  
signature-ca | {csr-cert {ocsp [on | off]}} file_name
```

Description

Permits downloading of certificate file(s) from files stored on a TFTP server.

Syntax Description

<i>ipaddress</i>	Specifies the IP address of the TFTP server.
ssl-cert	Specifies SSL/TLS certificate (default).
trusted-ca	Specifies CA certificates.
ocsp-signature-ca	Specifies signature CA files.
<i>file_name</i>	Specifies the name of the certificate file.
csr-cert	Specifies an SSL/TLS certificate signed through a Certificate Signing Request (CSR) generated by the switch. Trust chain verification is performed during download.
ocsp	Specifies using or not using Online Certificate Status Protocol (OCSP) for certificate checking.
on	Enables OCSP for SSL/TLS certificate signed through CSR generated by switch.
off	Disables OCSP for SSL/TLS certificate signed through CSR generated by switch (off).

Default

If no option is selected, SSL/TLS certificate (**ssl-cert**) is the default.

By default, OCSP is disabled.

Usage Guidelines

If the download operation is successful, any existing certificate is overwritten. For SSL/TLS certificates, after a successful download, the software attempts to match the public key in the certificate against the private key stored. If the private and public keys do not match, the switch displays a warning message similar to the following: Warning: The Private Key does not match with the Public Key in the certificate. This warning acts as a reminder to also download the private key.



Note

You can only download a certificate key in the *VR-Mgmt* virtual router.

Downloaded certificates and keys are not saved across switch reboots unless you save your current switch configuration. After you issue the `save` command, the downloaded certificate is stored in the configuration file and the private key is stored in the EEPROM.

You can purchase and obtain SSL certificates from Internet security vendors.

Remote IP Address Character Restrictions

This section provides information about the characters supported by the switch for remote IP addresses.

When specifying a remote IP address, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z).
- Numerals (0-9).
- Period (.).
- Colon (:).

When configuring an IP address for your network server, remember the requirements listed above.

Remote File Name Character Restrictions

This section provides information about the characters supported by the switch for remote file names.

When specifying a remote file name, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z).
- Numerals (0-9).
- Period (.).
- Dash (-).
- Underscore (_).
- Slash (/).

When naming a remote file, remember the requirements listed above.

Example

The following command downloads a certificate from a TFTP server with the IP address of 123.45.6.78:

```
# download ssl 123.45.6.78 certificate g0ethner1
```

The following command downloads a trusted-ca certificate:

```
# download ssl 10.120.89.79 certificate trusted-ca cacert.pem
```

The following command downloads an ocsp-signature-ca certificate:

```
# download ssl 10.120.89.79 certificate ocsp-signature-ca oscrcert.pem
```

History

This command was first available in the ExtremeXOS 11.2 and supported with the SSH module.

The **trusted-ca** and **ocsp-signature-ca** options were added in ExtremeXOS 22.1.

The **csr-cert** and **ocsp** were added in ExtremeXOS 31.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

download ssl privkey

```
download ssl ipaddress privkey key_file
```

Description

Permits downloading of a private key from files stored in a TFTP server.

Syntax Description

<i>ipaddress</i>	Specifies the IP address of the TFTP server.
<i>key_file</i>	Specifies the name of the private key file.

Default

N/A.

Usage Guidelines

If the operation is successful, the existing private key is overwritten.

After a successful download, a check is performed to find out whether the private key downloaded matches the public key stored in the certificate. If the private and public keys do not match, the switch displays a warning similar to the following: Warning: The Private Key does not match with the Public Key in the certificate. This warning acts as a reminder to also download the corresponding certificate.

The certificate and private key file should be in PEM format and generated using RSA as the cryptography algorithm.

Downloaded certificates and keys are not saved across switch reboots unless you save your current switch configuration. Once you issue the save command, the downloaded certificate is stored in the configuration file and the private key is stored in the EEPROM.

Remote IP Address Character Restrictions

This section provides information about the characters supported by the switch for remote IP addresses.

When specifying a remote IP address, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z).
- Numerals (0-9).
- Period (.).
- Colon (:).

When configuring an IP address for your network server, remember the requirements listed above.

Remote Filename Character Restrictions

This section provides information about the characters supported by the switch for remote filenames.

When specifying a remote filename, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z).
- Numerals (0-9).
- Period (.).
- Dash (-).
- Underscore (_).
- Slash (/).

When naming a remote file, remember the requirements listed above.

Example

The following command downloads a private key from a TFTP server with the IP address of 123.45.6.78:

```
download ssl 123.45.6.78 privkey t00Ts1e
```

History

This command was first available in the ExtremeXOS 11.2 and supported with the SSH module.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

edit policy

```
edit policy filename
```

Description

Edits a policy text file.

Syntax Description

<i>filename</i>	Specifies the filename of the policy text file.
-----------------	---

Default

N/A.

Usage Guidelines

This command edits policy text files that are on the switch. All policy files use “.pol” as the filename extension, so to edit the text file for the policy boundary use boundary.pol as the filename. If you specify the name of a file that does not exist, you will be informed and the file will be created.

This command spawns a VI-like editor to edit the named file. For information on using VI, if you are not familiar with it, do a web search for “VI editor basic information”, and you should find many resources. The following is only a short introduction to the editor.

Edit operates in one of two modes; command and input. When a file first opens, you are in the command mode. To write in the file, use the keyboard arrow keys to position your cursor within the file, then press one of the following keys to enter input mode:

- i - To insert text ahead of the initial cursor position.
- a - To append text after the initial cursor position.

To escape the input mode and return to the command mode, press the Escape key.

There are several commands that can be used from the command mode. The following are the most commonly used:

- dd - To delete the current line.
- yy - To copy the current line.
- p - To paste the line copied.
- :w - To write (save) the file.
- :q - To quit the file if no changes were made.
- :q! - To forcefully quit the file without saving changes.
- :wq - To write and quit the file.

Refresh Policy

After you have edited the text file for a policy that is currently active, you will need to refresh the policy if you want the changes to be reflected in the policy database. When you refresh the policy, the text file is read, the syntax is checked, the policy information is added to the policy manager database, and the policy then takes effect. Use the following command to refresh a policy:

```
refresh policy policy-name
```

If you just want to check to be sure the policy contains no syntax errors, use the following command:

```
check policypolicy-name{access-list}
```

Example

The following command allows you to begin editing the text file for the policy boundary:

```
edit policy boundary.pol
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

edit upm profile

```
edit upm profile profile-name
```

Description

Allows you to edit the specified profile.

Syntax Description

<i>profile-name</i>	Specifies the UPM profile to be edited.
---------------------	---

Default

N/A.

Usage Guidelines

Use the command to have VI-like editor features for editing the profile. Changes appear when you close the file for editing, not when you save it.

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the Universal Port feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

eject usb

```
eject usb-device
```

Description

Ensures that USB 2.0 storage device can be safely removed from the switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

After the switch writes to a USB 2.0 storage device, and before you can view the contents on the device, you must ensure it is safe to remove the device from the switch. Use this command to prepare the device for removal. After you run this command, you can manually remove the device.

If you have configured the [configure debug core-dumps](#) on page 393 command to write files to the device that you are trying to eject, you are reminded to select another location to write the debug files to:

Note: The destination of debug core dump is still configured to memorycard.
If a memory card will not be present, it is recommended to use "configure debug core-dumps" to change the core dump destination.

For more information about removing a USB 2.0 storage device, see the hardware documentation.

To access and read the data on the card, use a PC with appropriate hardware such as a compact flash reader/writer and follow the manufacturer's instructions to access the compact flash card and read the data.

Example

The following command prepares a compact flash card or USB 2.0 storage device to be removed from the switch:

```
# eject usb
```

History

This command was first available in ExtremeXOS 11.1.

Support for USB 2.0 storage devices was added in ExtremeXOS 12.5.3.

The **memorycard** was deprecated in ExtremeXOS 30.7.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

ELSE

ELSE



Note

This is a script statement and operates only in scripts when scripting is enabled with the following command: `enable cli scripting {permanent}`.

Description

Command block to be executed if the condition specified in the associated IF statement is not met.

Syntax Description

statements	Actions to be executed when the conditions specified in the associated IF statement are not met.
------------	--

Default

N/A.

Usage Guidelines

CLI scripting must be enabled before using this command.

This command must be preceded by `IF _expression THEN statements` and followed by `ENDIF`.

You can insert comments by using a number sign (#).

Example

The following example executes the `show switch` command if the value of the variable `x` is greater than 2, and execute the `show vlan` command otherwise:

```
IF ($x > 2) THEN
    show switch
```

```
ELSE
    show vlan
ENDIF
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable access-list permit to-cpu

```
enable access-list permit to-cpu
```

Description

Enables control packets to reach CPU, even if an ACL would deny them.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

This command allows control packets to reach the CPU, even if the packets match ACLs that would otherwise deny them. The control packets include STP and EAPS BPDUs, and ARP replies for the switch.

If this feature is disabled, these same packets will be denied if an ACL is applied that contains a matching entry that denies the packets. Contrary to expectations, when this feature is disabled, the packets will still be denied if there is a higher precedence entry that permits the packets.

To disable this feature, use the following command:

```
disable access-list permit to-cpu
```

Example

The following command enables STP BPDU packets to reach the switch CPU, despite any ACL:

```
enable access-list permit to-cpu
```

History

This command was first available in ExtremeXOS 11.3.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable access-list refresh blackhole

Enables blackholing of packets during *ACL* refresh.

```
enable access-list refresh blackhole
```

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

When access control lists (ACLs) are refreshed, this command provides that any packets arriving during the refresh will be blackholed. As the ACL is being refreshed, packets may arrive while the ACL is in an indeterminate state, and packets may be permitted that otherwise are dropped. This feature protects the switch during an ACL refresh.

To disable this feature, use the following command:

```
disable access-list refresh blackhole
```

Example

The following command enables dropping of packets during an ACL refresh:

```
enable access-list refresh blackhole
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable account

```
enable account [ all { admin | user | name } ]
```

Description

Enables the specified account locally.

Syntax Description

all	Specifies that all accounts, or all accounts of a certain type, will be enabled locally.
admin	Specifies that administrative privileged accounts will be enabled locally.
user	Specifies that user privileged accounts, including lawful intercept accounts, will be enabled locally.
<i>name</i>	Specifies the name of the account that will be enabled locally.

Default

Enabled.

Usage Guidelines

Enabling accounts affects the following northbound interfaces:

- Console
- TELNET
- SSH
- HTTP
- XML

Example

The following example enables all accounts locally:

```
enable account all
```

History

This command was first available in ExtremeXOS 15.7.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable avb

```
enable avb
```

Description

This command is a macro command that can be used to enable all AVB protocols globally on the switch. It is equivalent to issuing the following three commands:

```
enable mvrp
```

```
enable msrp
```

```
enable network-clock gptp
```

Syntax Description

avb	Audio Video Bridging.
------------	-----------------------

Default

Disabled.

Usage Guidelines

Use this command to enable all AVB protocols globally on the switch.



Note

AVB is not supported on 5720 VIM ports in Release 32.2.

Example

```
enable avb
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable avb ports

```
enable avb ports [port_list | all]
```

Description

This command is a macro command that can be used to enable all AVB protocols on a given set of ports. It is equivalent to issuing the following three commands:

```
enable mvrp ports
```

```
enable msrp ports
```

```
enable network-clock gtp ports
```

Syntax Description

avb	Audio Video Bridging.
<i>port_list</i>	Port list separated by a comma or "-".
all	All ports.

Default

Disabled.

Usage Guidelines

Use this command to enable all AVB protocols on the given ports.

Example

```
enable avb ports 1-5
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable bgp

```
enable bgp
```

Description

Enables *BGP*.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command enables the Border Gateway Protocol (BGP) on the router. Before invoking this command, the local AS number and BGP router ID must be configured.

Example

The following command enables BGP:

```
enable bgp
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

enable bgp advertise-inactive-route

```
enable bgp {address-family [ipv4-unicast | ipv4-multicast | ipv6-unicast | ipv6-multicast]} advertise-inactive-route
```

Description

Enables advertisement of BGP inactive routes, which are defined as those routes that are rated best by BGP and not best in the IP routing table.

Syntax Description

address-family	Specifies an IPv4 or IPv6 unicast or multicast address family.
-----------------------	--

Default

Disabled.

If no address family is specified, IPv4 unicast is the default address family.

Usage Guidelines

This command can be successfully executed only when BGP is globally disabled. It is best to enable this feature before you enable BGP ([enable bgp](#)). If BGP is enabled, you must disable BGP ([disable bgp](#)), enable this feature, and then enable BGP.

If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.

Example

The following command enables inactive route advertisement for IPv4 unicast traffic:

```
enable bgp address-family ipv4-unicast advertise-inactive-route
```

History

This command was first available in ExtremeXOS 12.1.3.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

enable bgp aggregation

```
enable bgp aggregation
```

Description

Enables *BGP* route aggregation.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

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:

```
configure bgp add aggregate-address {address-family [ipv4-unicast |  
ipv4-multicast | ipv6-unicast | ipv6-multicast]} ipaddress/masklength  
{as-match | as-set} {summary-only} {advertise-policy policy}  
{attribute-policy policy}
```

Example

The following command enables BGP route aggregation:

```
enable bgp aggregation
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

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

Disabled.

Usage Guidelines

MED is only used when comparing paths from the same AS, unless always-compare-med is enabled. When this command is issued, MEDs from different AS are used in comparing paths. A MED value of zero is treated as the lowest MED and therefore the most preferred route.

BGP must be disabled before you can change the configuration with this command.

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 ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

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` commands.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

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 AS-number : number
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the *BGP* feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

enable bgp export

For IPv4 and IPv6 routes:

```
enable bgp export route_type {{address-family} address_family} {export-policy policy-name}
```

For VPNv4 routes:

```
enable bgp export remote-vpn {{address-family} ipv4-unicast} {export-policy policy-name}
```

Description

For IPv4 and IPv6 routes, this command enables the export of routes learned from *BGP* peers to the specified protocol.

For VPNv4 routes, this command enables the exchange of routes between a BGP PE router and a CE router.

Syntax Description

bgp	For Layer 3 VPNs, this specifies that BGP routes learned from CE routers are to be exported to remote PE routers.
<i>route_type</i>	Specifies the BGP export route type. Valid <i>route_type</i> values are: blackhole; direct; isis; isis-level-1; isis-level-2; isis-level-1-external; isis-level-2-external; ospf; ospf-extern1; ospf-extern2; ospf-inter; ospf-intra; rip; static; ospfv3; ospfv3-extern1; ospfv3-extern2; ospfv3-inter; ospfv3-intra; ripng;
<i>address-family</i>	Valid <i>address_family</i> values are: ipv4-unicast; ipv4-multicast; ipv6-unicast; ipv6-multicast
remote-vpn	For Layer 3 VPNs, this specifies that BGP routes learned from remote PE routers are to be exported to the local VRF.
<i>policy-name</i>	Name of policy to be associated with network export. Policy can filter and/or change the route parameters.

Default

Disabled.

If no address family is specified for an IPv6 protocol, the default IPv6 unicast family applies; otherwise if no address family is specified, IPv4 unicast is the default.

Usage Guidelines

The exporting of routes between any two routing protocols is a discrete 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 a policy to associate BGP attributes including Community, NextHop, MED, Origin, and Local Preference with the routes. A policy can also be used to filter out exported routes.

Using the export command to redistribute routes complements the redistribution of routes using the `configure bgp add network` command. The `configure 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.



Note

For this command to execute, the specified protocol must support the specified address family. For example, the command fails if you specify OSPF and the IPv6 unicast address family. You can specify blackhole, direct, static, and IS-IS routes with IPv4 or IPv6 address families.

To export Layer 3 VPN routes to the CE peer in a VPN VRF, the source must be remote-vpn and destination address family must be ipv4-unicast.

Example

The following command enables BGP to export OSPF routes to other BGP routers:

```
enable bgp export ospf
```

The following command enables export of Layer 3 VPN Routes received from the PE Core in a VPN-VRF to its CE peers:

```
enable bgp export remote-vpn address-family ipv4-unicast
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

The blackhole option was added in ExtremeXOS 12.1.3.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Support for Layer 3 VPNs was added in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

enable bgp export vr

```
enable bgp export {vr} vr_name route_type {address-family} vpn4
    {export-policy policy_name}
```

Description

For IPv4 and IPv6 routes, this command enables the PE router to export and redistribute local VRF routes to remote PE routers through [BGP](#).

Syntax Description

vr	Specifies the source VPN VRF of the exported routes .
<i>vr_name</i>	Specifies the name of the source VPN VRF.
<i>route_type</i>	Specifies the source or origin of the route types to be exported to remote PE routers. Valid Types: blackhole, direct, and bgp, and static .
<i>address-family</i>	Specifies the address family for the exported routes. Valid types are vpn4.
export-policy vpn4	(Optional) The export policy can be specified when you enable bgp export. Specifies that routes from the VRF are exported as vpn4 routes over MPBGP.
<i>policy_name</i>	Name of export policy to be associated with export of VRF routes into BGP's VPN-IPv4 domain for advertisement to other PE routers.

Default

Disabled.

Usage Guidelines

This command enables a PE router to advertise learned routes from CE routers to remote PE routers in a Service Provider's backbone. Executing this command allows the PE router to convert VRF native IPv4 routes into VPN-IPv4 route,s and advertise to all remote PE BGP neighbors as VPN-IPv4 routes.

- This export command is applicable in Parent VR context only. If you execute it in a VRF context, an error message is returned.
- The source VPN VRF must be a child of the Parent VR.

- BGP need not be added to a VPN VRF to export routes from a VPN VRF.
- The direction of where the redistribution is targeted is implicit on the keywords used. Similarly `bgp` only applies to EBGp routes from CE exported as VPN routes, hence we use it only with address family `vpn4`. Other sources such as “static” and “direct” are redistributed both ways.
- Use `show vr parent_vr_name` to check routes exported from various VPN VRFs into the MBGP's VPN-IPv4 domain.
- Use `show vr vpn_vrf_name` to check routes exported from a VPN VRF into the MBGP's VPN-IPv4 domain.

Example

The following command enables BGP to advertise a `vpn4` route named “corp1_vpn_vrf”:

```
switch 19 # enable bgp export "corp1_vpn_vrf" bgp address-family vpn4
```

History

This command was first added in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

enable bgp export [static | direct] l2vpn-evpn

```
enable bgp export [static | direct] {address-family address_family}  
      l2vpn-evpn {vr vr_name}
```

Description

Exports direct, static, and BGP routes from a VRF into BGP, running on the specified VR, as EVPN routes to be advertised by BGP as Type 5 routes.

Syntax Description

bgp	Specifies showing the BGP configuration.
export	Specifies redistributing information from another routing protocol.
static	Specifies static routes.
direct	Specifies direct routes.
address-family	Specifies the address family.
<i>address_family</i>	Sets the address family type.
l2vpn-evpn	Specifies the L2VPN EVPN address family.

vr	Specifies the source VR.
<i>vr_name</i>	Designates the source VR name. Both VPN-VRFs and non-VPN-VRFs are supported.

Default

N/A.

Usage Guidelines

N/A.

Example

The following example exports static routes on VR "vr-a" as EVPN routes to be advertised by BGP as Type 5 routes:

```
# enable bgp export static l2vpn-evpn vr vr-a
```

History

This command was first available in ExtremeXOS 30.6.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

enable bgp fast-external-fallover

```
enable bgp fast-external-fallover
```

Description

Enables *BGP* fast external fallover functionality.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command enables the BGP fast external fallover on the router. This command applies to all directly-connected external BGP neighbors.

When BGP fast external fallover is enabled, the directly-connected EBGP neighbor session is immediately reset when the connecting link goes down.

If BGP fast external fallover is disabled, BGP waits until the default hold timer expires (3 keepalives) to reset the neighboring session. In addition, BGP might teardown the session somewhat earlier than hold timer expiry if BGP detects that the TCP session and it's directly connected link is broken (BGP detects this while sending or receiving data from TCP socket).

Example

The following command enables BGP fast external fallover:

```
enable bgp fast-external-fallover
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

enable bgp mpls-next-hop

```
enable bgp mpls-next-hop
```

Description

Enables IP forwarding over calculated MPLS LSPs to subnets learned via BGP.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command enables IP forwarding over calculated MPLS LSPs to subnets learned via BGP. (Calculated refers to an LSP that only reaches part of the way to the destination). IP forwarding over

MPLS LSPs must be enabled to forward over calculated LSPs. By default, IP forwarding over MPLS LSPs to subnets learned via BGP is disabled.

Example

The following command enables BGP's use of MPLS LSPs to reach BGP routes:

```
enable bgp mpls-next-hop
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

enable bgp multipath-relax

```
enable bgp multipath-relax
```

Description

Enables BGP multipath-relax feature, which modifies the definition of an equal cost BGP route.

Syntax Description

multipath-relax	Selects BGP multipath relax feature.
------------------------	--------------------------------------

Default

This feature is disabled by default.

Usage Guidelines

This feature modifies the definition of equal cost BGP routes as specified in *RFC-4271*. In particular, routes with the same AS-path length, but differing AS numbers in the path are not considered equal cost by default. However, with multipath-relax enabled, routes with the same AS-path length can have differing AS number values in the AS-path and still be considered equal cost.

BGP must be disabled (`disable bgp`) first to enable this feature.

Example

The following example enables BGP multipath-relax feature:

```
enable bgp multipath-relax
```

History

This command was first available in ExtremeXOS 22.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

enable bgp neighbor

```
enable bgp neighbor [remoteaddr | all]
```

Description

Enables the *BGP* session. The neighbor must be created before the BGP neighbor session can be enabled.

Syntax Description

<i>remoteaddr</i>	Specifies the IPv4 or IPv6 address of a BGP neighbor.
<i>all</i>	Specifies all IPv4 and IPv6 neighbors.

Default

Disabled.

Usage Guidelines

To create a new neighbor and add it to a BGP peer group, use the following command:

```
create bgp neighbor remoteaddr peer-group peer-group-name {multi-hop}
```

This command applies to the current VR or VRF context.

Example

The following command enables the BGP neighbor session:

```
enable bgp neighbor 192.168.1.17
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

enable bgp neighbor address-family l2vpn-evpn next-hop-unchanged

```
enable bgp {neighbor [remoteaddr | all]} [{address-family} l2vpn-evpn]
next-hop-unchanged
```

Description

Enables overriding the *BGP* specification behavior with respect to the next-hop of routes advertised to EBGp peers.

Syntax Description

bgp	Specifies BGP.
neighbor	Specifies BGP neighbor.
<i>remoteaddr</i>	Specifies BGP neighbor IP address.
all	Specifies all BGP neighbors.
address-family	Specifies address family.
l2vpn-evpn	Specifies L2VPN EVPN address-family type.
next-hop-unchanged	Enables preserving the BGP next-hop when routes are advertised to EBGp peers (default is disabled).

Default

Default is that next-hop-unchanged is disabled.

Usage Guidelines

This command enables overriding the specification behavior with respect to the next-hop of routes advertised to EBGp peers. Specifically, it maintains the BGP next-hop for routes advertised to EBGp peers instead of replacing the next-hop with either the outgoing interface IP address or the local loopback address.

When enabling the address family l2vpn-evpn in an EBGp configuration, the option **next-hop-unchanged** must also be enabled.

Example

The following example enables next-hop unchanged for BGP neighbor at 192.168.66.2:

```
# enable bgp neighbor 192.168.66.2 l2vpn-evpn next-hop-unchanged
```

History

This command was first available in ExtremeXOS 30.2.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

enable bgp neighbor capability

```
enable bgp neighbor [all | remoteaddr] capability [ipv4-unicast | ipv4-multicast | ipv6-unicast | ipv6-multicast | vpnv4 | route-refresh | ipv4-vxlan | l2vpn-evpn]
```

Description

This command enables multi protocol *BGP* (MBGP) and route-refresh capabilities for one or all BGP neighbors.

Syntax Description

all	Specifies that the configuration applies to all neighbors in the specified address family. If no address family is specified or if an IPv4 address is specified, the configuration applies to all IPv4 neighbors. If an IPv6 address family is specified, the configuration applies to all IPv6 neighbors.
<i>remoteaddr</i>	Specifies the IPv4 or IPv6 address of a BGP neighbor. The switch uses the IP address format to determine if the address is an IPv4 or IPv6 address.
ipv4-unicast	Specifies the IPv4 unicast address family.
ipv4-multicast	Specifies an IPv4 multicast address family.
ipv6-unicast	Specifies the IPv6 unicast address family.
ipv6-multicast	Specifies an IPv6 multicast address family.
vpnv4	Specifies VPN ipv4 unicast address family for a BGP neighbor. This is a required configuration for PE to PE BGP neighbor session. You must configure it before you enable a neighbor.
route-refresh	Specifies ROUTE-REFRESH message capabilities.
ipv4-vxlan	Specifies IPv4 VXLAN capability.
l2vpn-evpn	Specifies L2 VPN EVPN address family.

Default

The following capabilities are enabled by default for IPv4 peers: IPv4 unicast, IPv4 multicast, and route refresh.

The following capabilities are enabled by default for IPv6 peers: route refresh.

**Note**

For IPv4 peers, the IPv4 unicast and IPv4 multicast capabilities are enabled by default to support legacy peers that do not support MBGP. All other capabilities (except route-refresh) are disabled by default.

Usage Guidelines

When you change the capability configuration, you must enable the BGP neighbor before the configuration becomes active. If the BGP neighbor was enabled before the change, you must disable and enable the BGP neighbor. After the capabilities have been enabled, the BGP neighbor announces its capabilities to neighbors in an OPEN message.

When one or more address families are enabled, routes from the specified address families are updated, accepted, and installed. If more than one address family capability is enabled, or if the VPNv4 address family is enabled, the MBGP extension is automatically enabled. To disable MBGP, you must disable all enabled address families.

To support Layer 3 VPNs, you must enable the VPNv4 address family for all MBGP peers that will distribute VPNv4 routes across the service provider backbone. The VPNv4 address family must be enabled on the MPLS-enabled VR; it is not supported for BGP neighbors on the CE (VRF) side of the PE router.

Use the **vpnvp4** keyword for all PE to PE BGP neighbor sessions. This instructs BGP to negotiate the vpnvp4 address family in an open message with other PE routers. If this command is executed when a BGP neighbor session is established, it will take effect only after BGP session is reset. We recommend that you execute this command when a BGP neighbor is operationally down. Do not issue this command for a neighbor that is part of a VRF (PE – CE), or a warning message will be displayed.

**Note**

To inter-operate with Cisco routers for BGP graceful restart, you must enable the IPv4 unicast address capability.

This command applies to the current VR or VRF context.

**Note**

For an IPv6 peer, an IPv6 address family must be specified. From 21.1 ExtremeXOS allows IPV4 peering sessions to carry IPV6 routes and IPV6 peering sessions to carry IPV4 routes for the Unicast and Multicast sub-address families.

**Note**

You must enable a VPN IPv4 unicast address family for a BGP neighbor for a PE to PE BGP neighbor session before you enable the neighbor.

Example

The following command enables the route-refresh feature for all neighbors:

```
enable bgp neighbor all capability route-refresh
```

The following command enables the VPNv4 address family for a BGP neighbor:

```
virtual router corp1_vrf
enable bgp neighbor 192.168.96.235 capability vpnv4
```

The following command enables VXLAN capability for the BGP neighbor at 192.168.68.1:

```
enable bgp neighbor 192.168.68.1 capability ipv4-vxlan
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Support for Layer 3 VPNs was added in ExtremeXOS 15.3.

Support for IPv4 VXLAN was added in ExtremeXOS 22.3.

Support for L2 VPN EVPN address family was added in ExtremeXOS 30.2.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

enable bgp neighbor capability address-family vpnv4

```
enable bgp {neighbor} [all | remoteaddr] capability address-family vpnv4
type [community | ext-community] {[send | receive | both]}
```

Description

This command enables Outbound Route Filtering (ORF) for one or all [BGP](#) neighbors on a Layer 3 VPN.

Syntax Description

all	Specifies that the configuration applies to all neighbors in the specified address family. If no address family is specified or if an IPv4 address is specified, the configuration applies to all IPv4 neighbors.
<i>remoteaddr</i>	Specifies the IPv4 address of a BGP neighbor.
vpnv4	Specifies the VPNv4 address family for Layer 3 VPN support.
community	Enables neighbor capability for communities.
ext-community	Enables neighbor capability for extended communities.
send	Enables neighbor capability filter list send capability.

receive	Enables neighbor capability filter list receive capability.
both	Enables neighbor capability filter list send and receive capability.

Default

Disabled.

If the direction is not specified, the both option applies.

Usage Guidelines

Enter this command multiple times to configure the address family, type, and direction attributes.

By specifying the address-family, type and direction in multiple commands, you can better control the actual ORF capabilities sent to a peer. In the case where a particular address-family is explicitly disabled for a peering, the ORF capability configuration for that address-family is ignored and not sent.

ORF capabilities can only be enabled for IPv4 neighbors, and only for IPv4 address families. If configured for IPv6 neighbors or address-families the command is rejected with the following error message:

```
Outbound-route-filtering not supported for IPv6 neighbors
or
Outbound-route-filtering not supported for address family addr_family
```

Example

The following examples enables the neighbor capability feature for a Layer 3 VPN neighbor:

```
enable bgp neighbor 1.1.1.1 capability address-family vpnv4
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

enable bgp neighbor originate-default

```
enable bgp [{neighbor} remoteaddr | neighbor all] {address-family [ipv4-unicast | ipv4-multicast | ipv6-unicast | ipv6-multicast]} originate-default {policy policy-name}
```


Description

Enables the origination and advertisement of a default route to a single *BGP* neighbor or to all BGP neighbors.

Syntax Description

<i>remoteaddr</i>	Specifies the IPv4 or IPv6 address of a BGP neighbor.
all	Specifies that the configuration applies to all neighbors in the specified address family. If no address family is specified, the configuration applies to the IP Unicast family on all IPv4 peers. If an IPv4 address family is specified, the configuration applies to all IPv4 neighbors. If an IPv6 address family is specified, the configuration applies to all IPv6 neighbors.
address-family	Specifies an IPv4 or IPv6 unicast or multicast address family.
<i>policy-name</i>	Specifies a policy to be applied to the default route origination.

Default

Disabled. BGP does not automatically originate and advertise default routes to BGP neighbors.

If no address family is specified and an IPv4 address is detected, IPv4 unicast is the default address family.



Note

You must specify an IPv6 address family for an IPv6 peer, because an IPv6 peer does not support the default IPv4 unicast address family. Similarly, if you specify an IPv4 peer and an address family in the command, an IPv4 address family must be specified.

Usage Guidelines

This command can be successfully executed at any time, irrespective of whether local BGP or the remote BGP peer is enabled or disabled. The default route or routes are created regardless of whether or not there are matching entries in the IP route table.

When a BGP neighbor is added to a peer group, it does not inherit the default route origination configuration from the peer group. Also, default route origination for a neighbor and the associated peer group can be different.

If a policy is configured and specified in the command, a default route can be originated only if there is a route in the local BGP RIB that matches the policy's match rules. The default route's attribute can be modified using the same policy file by including statements in the set block of the policy.

If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.

Example

The following command enables the origination and advertisement of default routes for IPv4 unicast traffic for all BGP peer nodes:

```
enable bgp neighbor all originate-default
```

History

This command was first available in ExtremeXOS 12.3.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

enable bgp neighbor remove-private-AS-numbers

```
enable bgp neighbor [remoteaddr | 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

<i>remoteaddr</i>	Specifies the IPv4 or IPv6 address of a BGP neighbor.
all	Specifies all IPv4 and IPv6 neighbors.

Default

Disabled.

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.

This command applies to the current VR or VRF context.

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 192.168.1.17 remove-private-as-numbers
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

enable bgp neighbor soft-in-reset

```
enable bgp neighbor [all | remoteaddr] {address-family [ipv4-unicast |  
ipv4-multicast | ipv6-unicast | ipv6-multicast | vpn4]} soft-in-reset
```

Description

Enables the soft input reset feature.

Syntax Description

all	Specifies that the configuration applies to all neighbors in the specified address family. If no address family is specified, the configuration applies to the IP Unicast family on all IPv4 peers. If an IPv4 address family is specified, the configuration applies to all IPv4 neighbors. If an IPv6 address family is specified, the configuration applies to all IPv6 neighbors.
<i>remoteaddr</i>	Specifies the IPv4 or IPv6 address of a <i>BGP</i> neighbor.
ipv4-unicast	Specifies the IPv4 unicast address family.
ipv4-multicast	Specifies an IPv4 multicast address family.
ipv6-unicast	Specifies the IPv6 unicast address family.
ipv6-multicast	Specifies an IPv6 multicast address family.
vpn4	Specifies the VPNv4 address family for Layer 3 VPN support.

Default

Disabled.

If no address family is specified and an IPv4 address is detected, IPv4 unicast is the default address family.

Usage Guidelines

Before you can change the configuration with this command, you must disable BGP, and you must disable the corresponding BGP neighbor session using the following command:

```
disable bgp neighbor [remoteaddr | all]
```

To enable this feature on Layer 3 VPNs, you must do so in the context of the *MPLS*-enabled VR; this feature is not supported for BGP neighbors on the CE (VRF) side of the PE router.

This command applies to the current VR or VRF context.

If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.



Note

For an IPv6 peer, an IPv6 address family must be specified, because an IPv6 peer does not support IPv4 address families. If no address family is specified for an IPv6 peer, the default, the IPv4 unicast address family, applies and the command fails. Similarly an IPv4 peer only supports IPv4 address families and the command fails if an IPv6 address family is specified.

Example

The following command enables the soft recognition feature:

```
enable bgp neighbor 192.168.1.17 soft-in-reset
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Support for Layer 3 VPNs was added in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

enable bgp peer-group

```
enable bgp peer-group peer-group-name
```

Description

Enables a peer group and all the neighbors of a peer group.

Syntax Description

<i>peer-group-name</i>	Specifies a peer group.
------------------------	-------------------------

Default

Disabled.

Usage Guidelines

You can use [BGP](#) peer groups to group together up to 200512 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:

- source-interface
- out-nlri-filter
- out-aspath-filter
- out-route-map
- send-community
- next-hop-self

This command applies to the current VR or VRF context.

Example

The following command enables the BGP peer group outer and all its neighbors:

```
enable bgp peer-group outer
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

enable bgp peer-group capability

```
enable bgp peer-group peer-group-name capability {[address-family [ipv4-unicast | ipv4-multicast]]} type [community | ext-community | prefix]
{[send | receive | both]}
```

Description

This command enables ORF capabilities for a particular peer, peer-group, or all peers for one or all address-families and ORF types (for example, communities, extended communities and prefixes). The command specifies whether ORF capabilities are sent to the peer, and if they are honoured if received from the peer, or both.

Syntax Description

<i>peer-group-name</i>	Specifies a peer group.
address-family	Specifies outbound route filtering.
ipv4-unicast	Specifies an IPv4 unicast address family.
ipv4-multicast	Specifies an IPv4 multicast address family.
community	Enables ORF for communities.
ext-community	Enables ORF for extended communities.
prefix	Enables ORF for prefixes.
send	Enables ORF filter list send capability.
receive	Enables ORF filter list receive capability.
both	Enables ORF filter list send and receive capability.

Default

- ORF is disabled globally.
- ORF capabilities are assumed to be disabled by default for all neighbors.
- If address family is not specified, **ipv4-unicast** is assumed.
- If direction is not specified, **both** is assumed.



Note

prefix is not supported for vpnv4 address family.

The route refresh capability is enabled for IPv6 peer groups by default.

Usage Guidelines

By specifying the *address-family*, *type* and *direction* in multiple commands you can better control the ORF capabilities sent to a peer. In cases where a particular *address-family* is explicitly disabled for a peering, the ORF capability configuration for that address-family is ignored and not sent.

ORF capabilities can only be enabled for IPv4 neighbors and only for IPv4 address families. If configured for IPv6 neighbors or address-families the command is rejected with the following error message:

```
Outbound-route-filtering not supported for IPv6 neighbors, or Outbound-
route-filtering not supported for address family addr_family .
```

Example

The following command enables send only ORF capabilities for an ipv4 multicast peer group:

```
enable bgp peer-group capability orf address-family ipv4-multicast type community send
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the [BGP](#) feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

enable bgp peer-group capability

```
enable bgp peer-group peer-group-name capability [ipv4-unicast | ipv4-
multicast | ipv6-unicast | ipv6-multicast | vpn4 | route-refresh ]
```

Description

This command enables [BGP](#) Multiprotocol (MP) and route-refresh capabilities for a peer-group.

Syntax Description

<i>peer-group-name</i>	Specifies a peer group.
ipv4-unicast	Specifies the IPv4 unicast address family.
ipv4-multicast	Specifies an IPv4 multicast address family.
ipv6-unicast	Specifies the IPv6 unicast address family.
ipv6-multicast	Specifies an IPv6 multicast address family.
route-refresh	Specifies ROUTE-REFRESH message capabilities.
vpn4	Specifies the VPNv4 address family for Layer 3 VPN support.

Default

All capabilities are enabled for IPv4 peer groups by default.

The route refresh capability is enabled for IPv6 peer groups by default.

Usage Guidelines

This command enables BGP Multiprotocol or route-refresh capabilities for a peer group. When you change the capability configuration, you must enable the BGP peer group before the configuration becomes active. If the BGP peer group was enabled before the change, you must disable and enable the BGP peer group. After the capabilities have been enabled, the BGP peer announces its capabilities to neighbors in an OPEN message.

When one or more address families are enabled, routes from the specified address families are updated, accepted, and installed. If more than one address family capability is enabled, or if the VPNv4 address family is enabled, the MBGP extension is automatically enabled. To disable MBGP, you must disable all enabled address families.

A peer group can be configured for either IPv4 or IPv6 address families, but not both. Because a peer-group cannot support both IPv4 and IPv6 peers, the switch prevents the enabling of address families that are not compatible with peers that are already in the peer-group. Similarly if a particular address family is enabled for the peer-group, a peer that is incompatible with the existing peer-group configuration cannot be added to the group.

To support Layer 3 VPNs, you must enable the VPNv4 address family for all MBGP peers that will distribute VPNv4 routes across the service provider backbone. The VPNv4 address family must be enabled on the *MPLS*-enabled VR; it is not supported for BGP neighbors on the CE (VRF) side of the PE router.

This command applies to the current VR or VRF context.



Note

To inter-operate with Cisco routers for BGP graceful restart, you must enable IPv4 unicast address capability.

Example

The following command enables the route-refresh feature for the peer group outer:

```
enable bgp peer-group outer capability route-refresh
```

The following command enables the VPNv4 address family for a peer group:

```
enable bgp peer-group backbone capability vpnv4
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Support for Layer 3 VPNs was added in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

enable bgp peer-group capability address-family vpnv4

```
disable bgp peer-group peer-group-name capability address-family vpnv4
type [community | ext-community] {[send | receive | both]}
```

Description

This command disables peer-group capability for a peer group on a Layer 3 VPN.

Syntax Description

<i>remoteaddr</i>	Specifies the IPv4 address of a <i>BGP</i> neighbor.
vpn4	Specifies the VPNv4 address family for Layer 3 VPN support.
community	Disables peer-group capability for communities.
ext-community	Disables peer-group capability for extended communities.
send	Disables peer-group capability filter list send capability.
receive	Disables peer-group capability filter list receive capability.
both	Disables peer-group capability filter list send and receive capability.

Default

Disabled. If the direction is not specified, the **both** option applies.

Usage Guidelines

Enter this command multiple times to configure the address family, type, and direction attributes.

By specifying the address-family, type and direction in multiple commands, you can better control the actual ORF capabilities sent to a peer. In the case where a particular address-family is explicitly disabled for a peering, the ORF capability configuration for that address-family is ignored and not sent.

ORF capabilities can only be enabled for IPv4 neighbors, and only for IPv4 address families. If configured for IPv6 neighbors or address-families the command is rejected with the following error message:

```
Outbound-route-filtering not supported for IPv6 neighbors, orOutbound-
route-filtering not supported for address family addr_family .
```

The following command disables the peer-group capability feature for a Layer 3 VPN peer group:

```
disable bgp peer-group vpn capability address-family vpnv4
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

enable bgp peer-group next-hop-unchanged

```
enable bgp peer-group peer-group-name l2vpn-evpn next-hop-unchanged
```

Description

Enables a peer group and with respect to the next-hop of routes advertised to EBGp peers.

Syntax Description

<i>peer-group-name</i>	Specifies a peer group.
l2vpn-evpn	Specifies L2VPN EVPN address-family type.
next-hop-unchanged	Enables preserving the BGP next-hop when routes are advertised to EBGp peers (default is disabled).

Default

Default is that next-hop-unchanged is disabled.

Usage Guidelines

This command enables overriding the specification behavior with respect to the next-hop of routes advertised to EBGp peers. Specifically, it maintains the BGP next-hop for routes advertised to EBGp peers.

When enabling the address family l2vpn-evpn in an EBGp configuration, the option **next-hop-unchanged** must also be enabled.

Example

The following command enables next-hop unchanged for the BGP peer group pg2 :

```
enable bgp peer-group pg2 l2vpn-evpn next-hop-unchanged
```

History

This command was first available in ExtremeXOS 31.4.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

enable bgp peer-group originate-default

```
enable bgp {peer-group} peer-group-name {address-family [ipv4-unicast  
| ipv4-multicast | ipv6-unicast | ipv6-multicast] } originate-default  
{policy policy_name}
```

Description

Enables the origination and advertisement of default routes to all *BGP* neighbors in the specified peer group.

Syntax Description

peer-group <i>peer-group-name</i>	Specifies the BGP peer group for which the default routes are originated and advertised.
address-family	Specifies an IPv4 or IPv6 unicast or multicast address family.
<i>policy_name</i>	Specifies a policy to be applied to the default routes during origination.

Default

Disabled. BGP does not automatically originate and advertise default routes to BGP neighbors.

Usage Guidelines

This command can be successfully executed at any time, irrespective of whether local BGP or the remote BGP peers are enabled or disabled. The default routes are created regardless of whether or not there are matching entries in the IGP route table.

When a BGP neighbor is added to a peer group, it does not inherit the default route origination configuration from the peer group. Also, default route origination for a neighbor and the associated peer group can be different.

If a policy is configured and specified in the command, a default route can be originated only if there is a route in the local BGP RIB that matches the policy's match rules. The default route's attribute can be modified using the same policy file by including statements in the set block of the policy.

If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.



Note

If the specified peer group contains IPv6 peers, it is an IPv6 peer group and you must specify an IPv6 address-family. When the specified peer group is an IPv6 peer group, this command fails if no address family is specified or if an IPv4 address-family is specified. This command also fails if an IPv6 address family is specified for an IPv4 peer-group.

Example

The following command enables the origination and advertisement of default routes for IPv4 unicast traffic for all nodes in the test BGP peer group:

```
enable bgp peer-group test originate-default
```

History

This command was first available in ExtremeXOS 12.2.2.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

enable bgp peer-group remove-private-AS-numbers

```
enable bgp peer-group peer-group-name remove-private-AS-numbers
```

Description

Enables the removal of private autonomous system (AS) numbers from the AS_Path attribute of outbound updates.

Syntax Description

<i>peer-group-name</i>	Specifies a peer group.
------------------------	-------------------------

Default

Disabled.

Usage Guidelines

This command applies to the current VR or VRF context.

Example

The following command enables the *BGP* peer group outer from removing private AS numbers:

```
enable bgp peer-group outer remove-private-AS-numbers
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

enable bgp peer-group soft-in-reset

```
enable bgp peer-group peer-group-name {address-family [ipv4-unicast |  
ipv4-multicast | ipv6-unicast | ipv6-multicast | vpn4] }soft-in-reset
```

Description

Enables the soft input reset feature.

Syntax Description

<i>peer-group-name</i>	Specifies a peer group.
ipv4-unicast	Specifies the IPv4 unicast address family.
ipv4-multicast	Specifies an IPv4 multicast address family.
ipv6-unicast	Specifies the IPv6 unicast address family.
ipv6-multicast	Specifies an IPv6 multicast address family.
vpn4	Specifies the VPNv4 address family for Layer 3 VPN support.

Default

Disabled.

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

Disabling the soft input reset feature can potentially limit the amount of system memory consumed by the RIB-in.

After you enter this command, the switch automatically disables and enables all neighbors in the peer group before the change takes effect.

This command applies to the current VR or VRF context.

If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.



Note

If the specified peer group contains IPv6 peers, it is an IPv6 peer group and you must specify an IPv6 address-family. When the specified peer group is an IPv6 peer group, this command defaults to IPv4 unicast if no address family is specified. This command fails if an IPv6 address family is specified for an IPv4 peer-group.

Example

The following command enables the soft input reset feature:

```
enable bgp peer-group outer soft-in-reset
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 [BGP](#).

Support for Layer 3 VPNs was added in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

enable bootp vlan

```
enable bootp {ipv4} | dhcp {ipv4 | ipv6} ] vlan [vlan | 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

bootp	Enable BOOTP client.
ipv4	IPv4 client.
dhcp	Enable <u>DHCP</u> client.
ipv6	IPv6 client.
<i>vlan</i>	Specifies a VLAN name.
all	Specifies all VLANs.

Default

Disabled.

Usage Guidelines

If IPv4/IPv6 keyword is not specified , ipv4 would be taken as default for the mentioned VLAN.

Example

The following example enables the generation and processing of BOOTP packets on a VLAN named "accounting":

```
enable bootp vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

The **ipv4** and **ipv6** keywords were added in ExtremeXOS 15.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable bootprelay ipv6

```
enable bootprelay {ipv4 | ipv6} {vlan vlan_name} | {vr vr_name} | all  
                  {vr vr_name}
```

Description

Enables BOOTP Relay v6. This can be done across the VR or on a per VLAN basis.

Syntax Description

bootprelay	BOOTP Relay service.
ipv4	DHCPv4 BOOTP Relay service.
ipv6	DHCPv6 BOOTP Relay service.
<i>vlan_name</i>	Specifies a VLAN name
vr	Uses a specific virtual router name.
<i>vr_name</i>	Specifies a virtual router name.
all	Enables all VLANs.

Default

IPv4.

Usage Guidelines

Use this command to enable BOOTP Relay across the VR or on a per VLAN basis.

Example

The following example displays IPv6 bootprelay information:

```
* switch # show bootprelay ipv6
BOOTP Relay: DHCPv6 BOOTP Relay enabled on virtual router "VR-Default"
  BOOTP Relay Servers : 2001:0db8:85a3:0000:0000:8a2e:0370:7334
                        2001:0db8:85a3:0000:0000:8a2e:0370:7335
                        2001:0db8:85a3:0000:0000:8a2e:0370:7336
                        2001:0db8:85a3:0000:0000:8a2e:0370:7337

VLAN "Default":
  BOOTP Relay          : Disabled
VLAN "v1":
  BOOTP Relay          : Enabled
  BOOTP Relay Servers  : 2001:0db8:85a3:0000:0000:8a2e:0370:7338
  Interface ID         : v1-12
  Remote ID            : v1_remId
  Prefix Snooping      : Disabled
VLAN "v2":
  BOOTP Relay          : Enabled
  BOOTP Relay Servers  : 2001:0db8:85a3:0000:0000:8a2e:0370:7339
  Interface ID         : 100 (Default)
  Remote ID            : 00:04:96:52:A7:1B (Default)
  Prefix Snooping      : Disabled
```

History

This command was first available in ExtremeXOS 15.5.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

enable bootrelay

```
enable bootrelay [{vlan} [vlan_name] | [{vr} vr_name] | all [{vr}  
vr_name]}
```

Description

Enables the BOOTP Relay function on one or all VLANs for the specified VR or VRF.

Syntax Description

<i>vlan_name</i>	Specifies a single VLAN on which to enable the BOOTP Relay feature.
<i>vr_name</i>	Specifies a single VR or VRF on which to enable the BOOTP Relay feature. If not specified, VR or VRF of current command context is used.
all	Specifies that BOOTP Relay is to be enabled for all VLANs on the specified VR or VRF.

Default

The BOOTP Relay function is disabled on all VLANs and VRs.

If not specified, VR of current command context is used.

Usage Guidelines

Because VLAN names are unique on the switch, you can specify only a VLAN name (and omit the VR name) to enable BOOTP Relay on a particular VLAN. When you enable BOOTP Relay on a VR or VRF, BOOTP Relay is enabled on all VLANs for that VR. If you enter the command without specifying a VLAN or a VR, the functionality is enabled for all VLANs in the current VR context.



Note

If DHCP/BOOTP Relay is enabled on a per VLAN basis, make sure it is enabled on both the client-side and server-side VLANs.

Example

The following example enables the forwarding of BOOTP requests for all VLANs in the current VR context:

```
enable bootrelay
```

You can use either of the following commands to enable the forwarding of BOOTP requests for VLAN client1:

```
enable bootprelay "client1"  
enable bootprelay vlan "client1"
```

You can use any one of the following commands to enable the forwarding of BOOTP requests for all VLANs on VR zone3:

```
enable bootprelay zone3  
enable bootprelay vr zone3  
enable bootprelay all zone3  
enable bootprelay all vr zone3
```

History

This command was first available in ExtremeXOS 10.1.

The capability to enable BOOTP Relay on a VLAN was added in ExtremeXOS 12.4.2.

The capability to enable BOOTP Relay on VPN-VRF is added in ExtremeXOS 22.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable cdp ports

```
enable cdp ports [port_list | all]
```

Description

Enables Cisco Discovery Protocol (CDP) on a port.

Syntax Description

<i>port_list</i>	Specifies the list of ports to enable CDP on.
all	Specifies that you enable CDP on all ports.

Default

Enabled.

Usage Guidelines

Example

The following command enables CDP on all ports on the switch:

```
enable cdp ports all
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable cfm segment frame-delay measurement

```
enable cfm segment frame-delay measurement segment_name { mep mep_id }  
[continuous | count value]
```

Description

Triggers DMM frame transmission.

Syntax Description

<i>segment_name</i>	An alpha numeric string identifying the segment name.
mep	Specifies the maintenance association End Point that helps trigger a particular MEP level session on that segment.
<i>mep_id</i>	Specifies the MEP-ID. The range is 1-8191. The default is all MEPs on the segment.
continuous	Specifies that frames are to be sent continuously until stopped.
count	Specifies that a number of frames are to be sent.
<i>value</i>	Specifies the number of frames to send. The range is 1 to 4294967295.

Default

N/A.

Usage Guidelines

Use this command to trigger DMM frames at the specified transmit interval configured using the command `configure cfm segment transmit-interval`.

Continuous transmission continues until it is stopped with the command `disable cfm segment frame-delay measurement` or `delete cfm segment`.



Note

If you try to trigger the DMM frames for a segment that is not completely configured, the frames are not transmitted for that segment, and an error message is displayed on the console.

Example

The following example triggers continuous frame transmission on the CFM segment segment-first:

```
enable cfm frame-delay measurement segment-first continuous
```

History

This command was first available in ExtremeXOS 12.3.

The **mep** keyword was added in ExtremeXOS 15.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable cfm segment frame-loss measurement mep

If the user specifies the mode as continuous, the LMM transmission will continue till it is stopped by the user.

```
enable cfm segment frame-loss measurement segment_name mep mep_id
[continuous | count frames]
```

Description

This command is used to trigger LMM frames at the configured transmit-interval.

Syntax Description

<i>segment_name</i>	An alpha numeric string identifying the segment name.
continuous	Specifies that frames are to be sent continuously until stopped.
count	Specifies that a number of frames are to be sent.
<i>frames</i>	Specifies the number of frames to send. The range is 1 to 4294967295.

Default

N/A.

Usage Guidelines

This command is used to trigger LMM frames at the configured transmit-interval. If the user specifies the mode as continuous, the LMM transmission will continue till it is stopped by the user.



Note

If the user tries to trigger the LMM frames for a segment which is not completely configured, the frames will not be transmitted for that segment, and an error message will be thrown.

Example

```
enable cfm segment cs2 frame-loss measurement mep 3 count 10
enable cfm segment cs2 frame-loss measurement mep 3 continuous
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable clear-flow

```
enable clear-flow
```

Description

Enable the CLEAR-Flow agent.

Syntax Description

This command has no arguments or variables.

Default

CLEAR-Flow is disabled by default.

Usage Guidelines

When the CLEAR-Flow agent is enabled, sampling begins and actions are taken based on the CLEAR-Flow rules that are configured on the switch.

Example

The following example enables CLEAR-Flow on the switch:

```
# enable clear-flow
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, 5720 series switches.

enable cli history expansion

```
enable cli history expansion {session | permanent}
```

Description

Performs command line history expansion similar to the Linux shell.

Syntax Description

cli	Command line interface settings.
history	Command history settings.
expansion	Substitute occurrences of '!n:w' with the corresponding line 'n' and word 'w+1' from command history (default disabled).
session	Configures history expansion for this CLI session only (default).
permanent	Configures history expansion for this CLI session, and all future sessions.

Default

CLI history expansion is disabled by default. If not specified when enabling, CLI history expansion is enabled for the current session only.

Usage Guidelines

The history expansion character '!' can be used to specify command from the history that is substituted into the command line. All occurrences of the form "!"n:w" in the command are replaced with the w'th word from the n'th line in the command history. Specification of the word is optional.

If you enable CLI history expansion, and then try to reference a history that does not exist, the following error appears:

```
# show !58:1 Error: History event not found. If you were not attempting a history
expansion using the format '!n:w', and believe the command to be valid, please retry the
command after 'disable cli history expansion'.
```

To view the status of CLI history expansion on the switch, use the `show management` command.

Example

The following command enables CLI history expansion for this session and all future sessions:

```
enable cli history expansion permanent
```

History

This command was first available in ExtremeXOS 32.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable cli prompting

```
enable cli prompting
```

Description

Enables CLI prompting for the session.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

Use this command to enable CLI prompting from a disabled state.

To view the status of CLI prompting on the switch, use the [show management](#) command.

Example

The following command enables prompting:

```
enable cli prompting
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable cli refresh

```
enable cli refresh {session | permanent}
```

Description

This command allows you to configure the default auto refresh behavior. The auto refresh behavior is used for some `show` commands.

Syntax Description

session	Use refresh setting for this CLI session only.
permanent	Use refresh setting for this CLI session, and all future sessions (default).

Default

Permanent.

Usage Guidelines

The auto refresh behavior is used for some ‘show’ commands. You must use the `disable cli refresh` command to disable the show command auto refresh or add the no-refresh option to the individual command. For example:

- `show ports config` - will display and refresh the first <n> ports of a switch until the **[ESC]** key is pressed.
- `disable cli refresh`
- `show ports config` - will act as if `show ports config no-refresh` was entered and page through all ports

Since the default for the session may be set to `disable cli refresh` the commands that take a **no-refresh** option now allow for the alternate **refresh** case if the user wants to selectively enable a refreshed display.

The **permanent** option is only valid for admin level users.

Example

The following sample output displays the CLI refresh information.

```
# show management
CLI idle timeout           : Enabled (20 minutes)
CLI max number of login attempts : 3
CLI max number of sessions : 8
CLI paging                 : Enabled (this session only)
CLI space-completion       : Disabled (this session only)
CLI configuration logging   : Disabled
CLI scripting              : Disabled (this session only)
CLI scripting error mode    : Ignore-Error (this session only)
CLI persistent mode        : Persistent (this session only)
```



```
CLI prompting          : Enabled (this session only)
CLI refresh            : Enabled (this session only)
Telnet access          : Enabled (tcp port 23 vr all)
                       : Access Profile : not set
SSH Access             : ssh module not loaded.
Web access             : Enabled (tcp port 80)
                       : Access Profile : not set
```

History

This command was first available in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable cli scripting

```
enable cli scripting {permanent}
```

Description

Enables the use of CLI scripting commands. When used without the permanent option, it enables the CLI scripting commands for the current session and is a per session setting. The permanent option enables the CLI scripting commands for new sessions only and is saved across switch reboots.

Syntax Description

permanent	Enables the CLI scripting commands for new sessions only; this setting is saved across switch reboots.
------------------	--

Default

The CLI scripting commands are disabled by default.

Usage Guidelines

You must enable the CLI scripting commands on the switch to use the scripting keywords in the script, and before you can configure or execute a script.



Note

CLI scripting commands cannot be enabled when CLI space auto completion is enabled with the `enable cli space-completion` command.

Example

The following command enables the CLI scripting commands for the current session:

```
enable cli scripting
```

History

This command was first available in ExtremeXOS 11.6.

The permanent option was added in ExtremeXOS 12.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable cli scripting output

```
enable cli scripting output
```

Description

Enables the display of CLI commands and responses during script operation.

Default

During interactive script sessions: CLI scripting output enabled.

During load script command operation: CLI scripting output disabled.

Usage Guidelines

When the CLI scripting output is enabled, all script commands and responses are displayed.

When the `load script filename {arg1} {arg2} ... {arg9}` command is entered, the software disables CLI scripting output until the script is complete, and then CLI scripting output is enabled. Use the `enable cli scripting output` and `disable cli scripting output` commands to control what a script displays when you are troubleshooting.

Example

The following command enables CLI scripting output for the current session or until the `disable cli scripting output` command is entered:

```
enable cli scripting output
```

History

This command was first available in ExtremeXOS 12.1.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable cli space-completion

```
enable cli space-completion
```

Description

Enables the ExtremeXOS feature that completes a command automatically with the spacebar. The [Tab] key can also be used for auto-completion.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

CLI space auto completion cannot be enabled while CLI scripting is enabled with the `enable cli scripting` command.

Example

The following command enables using the spacebar to automatically complete a command:

```
enable cli space-completion
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

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

Disabled.

Usage Guidelines

ExtremeXOS 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.

To view the status of configuration logging on the switch, use the `show management` command. This 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 ExtremeXOS 11.0.

The `cli-config-logging` keyword was split into `cli config-logging` in ExtremeXOS 30.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable cli-config-logging expansion

```
enable cli-config-logging expansion
```

Description

When CLI logging is enabled, enables showing fully expanded commands, rather than abbreviations, in the log.

Syntax Description

<code>expansion</code>	Enables command expansion in logs.
-------------------------------	------------------------------------

Default

Expansion is disabled by default.

Usage Guidelines

When CLI logging is enabled (see [enable cli config-logging](#) on page 2095), this command enables showing fully expanded commands, rather than abbreviations, in the log.

For example, with command expansion enabled, a command entered in abbreviated format, such as

```
config por 33 auto of spee 10000 duplex ful
```

appears in the log as:

```
configure ports 33 auto off speed 10000 duplex full
```

Whereas, if command expansion is turned off, the command appears in the log in the exact format as it was typed into the command line.

To see the status of command expansion, use [show management](#) on page 2854.

Example

The following example turns on command expansion:

```
enable cli-config-logging expansion
```

History

This command was first available in ExtremeXOS 22.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable cli paging

```
enable cli paging {session | permanent}
```

Description

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

Syntax Description

session	Enables viewing output of commands one screenful at a time for the current user session only (default).
permanent	Enables viewing output of commands one screenful at a time permanently (setting persists after rebooting).

Default

Enabled per session.

Usage Guidelines

The command line interface (CLI) is designed for use in a VT100 environment.

Most show command output pauses 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 cli paging permanently (setting persists across reboots) and does not allow the display to print continuously to the screen:

```
enable cli paging permanent
```

History

This command was first available in ExtremeXOS 10.1.

The **session** and **permanent** options were added in ExtremeXOS 22.5.

The **clipaging** option was split into two keywords in ExtremeXOS 30.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable cpu-monitoring

```
enable cpu-monitoring {interval seconds} {threshold percent}
```

Description

Enables CPU monitoring on the switch.

Syntax Description

<i>seconds</i>	Specifies the monitoring interval, in seconds. The default is 5 seconds, and the range is 5 to 60 seconds.
threshold	Specifies the CPU threshold value. CPU usage is measured in percentages. The default is 90%, and the range is 0% to 100%.

Default

CPU monitoring is enabled and occurs every 5 seconds. The default CPU threshold value is 90%.

Usage Guidelines

CPU monitoring allows you to monitor the CPU utilization and history for all of the processes running on the switch. By viewing this history on a regular basis, you can see trends emerging and identify processes with peak utilization. Monitoring the workload of the CPU allows you to troubleshoot and identify suspect processes before they become a problem.

To specify the frequency of CPU monitoring, use the interval keyword. We recommend the default setting for most network environments.

CPU usage is measured in percentages. By default, the CPU threshold value is 90%. When CPU utilization of a process exceeds 90% of the regular operating basis, the switch logs an error message specifying the process name and the current CPU utilization for the process. To modify the CPU threshold level, use the threshold keyword. The range is 0% to 100%.

Example

The following command enables CPU monitoring every 30 seconds:

```
enable cpu-monitoring interval 30
```

History

This command was first available in ExtremeXOS 11.2.

The default values shown began in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable dhcp ports vlan

```
enable dhcp ports port_list vlan vlan_name
```

Description

Enables *DHCP* on a specified port in a *VLAN*.

Syntax Description

<i>port_list</i>	Specifies the ports for which DHCP should be enabled.
<i>vlan_name</i>	Specifies the VLAN on whose ports DHCP should be enabled.

Default

Disabled.

Usage Guidelines

None.

Example

The following command enables DHCP for port 5:9 in VLAN corp:

```
enable dhcp ports 5:9 vlan corp
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable dhcp vlan

enable dhcp [ipv4 | ipv6] [vlan_name | **all**]

Description

Enables the generation and processing of *DHCP* packets on a *VLAN* to obtain an IP address for the VLAN from a DHCP server.

Syntax Description

<i>vlan_name</i>	Specifies a VLAN name.
all	Specifies all VLANs.

Default

If the IPv4/IPv6 keyword is not specified, IPv4 is taken as default for the mentioned VLAN | all.

Usage Guidelines

None.

Example

The following command enables the generation and processing of DHCP packets on a VLAN named accounting:

```
enable dhcp vlan accounting
enable dhcp ipv6 vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

This command was modified in ExtremeXOS 15.6 to include the **ipv4** and **ipv6** keywords.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable diffserv examination ports

```
enable diffserv examination ports [port_list | all]
```

Description

Enables the DiffServ field of an IP packet to be examined in order to select a [QoS](#) profile.

Syntax Description

<i>port_list</i>	Specifies a list of ports or slots and ports to which the parameters apply.
all	Specifies that DiffServ examination is enabled for all ports.

Default

Disabled.

Usage Guidelines

The Diffserv examination feature is disabled by default.

Example

The following command enables DiffServ examination on selected ports:

```
enable diffserv examination ports 1:1,5:5,6:2
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable diffserv replacement ports

```
enable diffserv replacement ports [port_list | all] [{qosprofile}
qosprofile}
```

Description

Enables the DiffServ code point to be overwritten in IP packets transmitted by the switch.

Syntax Description

<i>port_list</i>	Specifies a list of ingress ports or slots and ports on which to enable Diffserv replacement.
all	Specifies that DiffServ replacement should be enabled for all ports.
qosprofile	Enables DiffServ replacement on a QoS profile. Note: DiffServ replacement will be enabled for all QoS profiles if this option is not specified.
<i>qosprofile</i>	Specifies the QoS profile number.

Default

N/A.

Usage Guidelines

The Diffserv replacement feature functions for IPv4 and IPv6 traffic and is disabled by default.



Note

The port in this command is the ingress port.

This command affects only that traffic in traffic groupings based on explicit packet class of service information and physical/logical configuration.

Example

The following example enables DiffServ replacement on specified ports:

```
enable diffserv replacement ports 5:3,5:5,6:2
```

History

This command was first available in ExtremeXOS 11.0.

The **gosprofile** keyword and *gosprofile* variable were added in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable dns cache

```
enable dns cache [{vlan} vlan_name | {vr} vr_name]
```

Description

Enables the Domain Name System (DNS) cache on a virtual router (VR) or VLAN.

Syntax Description

dns	Domain name system.
cache	Specifies enabling the DNS cache.
vlan	Specifies enabling DNS cache on a VLAN.
<i>vlan_name</i>	Specifies the VLAN name.
vr	Specifies enabling DNS cache on a VR.
<i>vr_name</i>	Specifies the VR name. If not specified, VR of current command context is used.

Default

If no VR name is specified, the VR of the current command context is used.

Usage Guidelines

To view the DNS cache configuration, use the command `show dns cache configuration [{vlan} vlan_name | {vr} vr_name]`

Example

The following example enables DNS cache on VLAN "VLAN1":

```
# enable dns cache vlan VLAN1
```

History

This command was first available in ExtremeXOS 30.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable dns cache dnssec

```
enable dns cache {dnssec}
```

Description

Validates DNS replies and cache data for DNSSEC (Domain Name System Security Extensions).

Syntax Description

dnssec	Validate DNS replies and cache data for DNSSEC. Default is disabled.
---------------	--

Default

By default, DNSSEC is disabled.

Usage Guidelines

You cannot enable DNSSEC if DNS cache is enabled.

Example

The following example enables DNSSEC:

```
# enable dns cache dnssec
```

History

This command was first available in ExtremeXOS 31.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable dns cache analytics

```
enable dns cache analytics {{vr} vr_name}
```

Description

Enables Domain Name System (DNS) analytics.

Syntax Description

dns	Domain Name System.
cache	Specifies DNS cache.
analytics	Specifies enabling DNS cache analytics. Analytics provides more insight into DNS queries when DNS cache is enabled. Default is disabled.
vr	Specifies enabling DNS analytics on a VR.
<i>vr_name</i>	Specifies the VR name. If not specified, the VR of the current command context is used.

Default

DNS analytics is disabled by default.

Usage Guidelines

To disable DNS analytics, use the command `disable dns cache analytics {{vr} vr_name}`.

Example

The following example enables DNS analytics on VR "vr1":

```
# enable dns cache analytics vr vr1
```

History

This command was first available in ExtremeXOS 30.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable dos-protect simulated

```
enable dos-protect simulated
```

Description

Enables simulated denial of service protection.

Syntax Description

This command has no arguments or variables.

Default

The default is disabled.

Usage Guidelines

If simulated denial of service is enabled, no ACLs are created. This mode is useful to gather information about normal traffic levels on the switch. This will assist in configuring denial of service protection so that legitimate traffic is not blocked.

Example

The following command enables simulated denial of service protection:

```
enable dos-protect simulated
```

History

This command was first available in ExtremeXOS 11.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable dos-protect

```
enable dos-protect
```

Description

Enables denial of service protection.

Syntax Description

This command has no arguments or variables.

Default

The default is disabled.

Usage Guidelines

None.

Example

The following command enables denial of service protection.

```
enable dos-protect
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable dot1p examination inner-tag port

```
enable dot1p examination inner-tag port [all | port_list]
```

Description

Used with VMANs, and instructs the switch to examine the 802.1p value of the inner tag, or header of the original packet, to determine the correct egress queue on the egress port.

Syntax Description

all	Specifies all ports.
<i>port_list</i>	Specifies a list of ports or slots and ports.

Default

Disabled.

Usage Guidelines

Use this command to instruct the system to refer to the 802.1p value contained in the inner, or original, tag when assigning the packet to an egress queue at the egress port of the VMAN.



Note

For information about configuring and displaying the current 802.1p and DiffServ configuration for the inner, or original header, 802.1p value, see the *Quality of Service* section in the [Switch Engine 32.3 User Guide](#).

Example

The following example puts the packets in the egress queue of the VMAN egress port according to the 802.1p value on the inner tag:

```
enable dot1p examination inner-tag port 3:2
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable dot1p examination ports

```
enable dot1p examination ports [port_list | all]
```

Description

Enables egress QoS profile selection based on the 802.1p bits in the incoming frame.

Syntax Description

<i>port_list</i>	Specifies a list of ports on which to enable the dot1p examination feature.
all	Specifies that dot1p examination should be enabled for all ports.

Default

Enabled.

Usage Guidelines

To increase available ACLs, you can disable the 802.1p examination feature if you are not running QoS or are running QoS using DiffServ. See the [Switch Engine 32.3 User Guide](#) for information on ACL limitations on these platforms.

Use this command to re-enable the 802.1p examination feature.

As part of the COS global status enable action, COS will automatically enable dot1p examination on all ports. An internal status will track this event. The disable dot1p examination command will print an additional warning message in the event that COS was configured via SNMP. If the COS global status is disabled via SNMP, the internal status will be cleared and the additional warning message will not be displayed.

Example

The following command enables dot1p examination on ports 1 to 5:

```
enable dot1p examination ports 1-5
```

History

This command was available in ExtremeXOS 11.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable dot1p replacement ports

```
enable dot1p replacement ports [port_list | all] [{qosprofile}  
  qosprofile}
```

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

<i>port_list</i>	Specifies a list of ports or slots and ports.
all	Specifies that dot1p replacement should be enabled for all ports.
qosprofile	Enables dot1p on a QoS profile.
<i>qosprofile</i>	Specifies the QoS profile number.

Default

N/A.

Usage Guidelines

The dot1p replacement feature is disabled by default.

By default, 802.1p priority information is not replaced or manipulated, and the information observed on ingress is preserved when transmitting the packet.



Note

The port in this command is the ingress port.

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.



Note

This command affects only that traffic in traffic groupings based on explicit packet class of service information and physical/logical configuration.

Beginning with ExtremeXOS version 11.4 on the 1 Gigabit Ethernet ports, 802.1p replacement always happens when you configure the DiffServ traffic grouping.



Note

Enabling dot1p replacement on all ports may take some time to complete.

Example

The following example enables dot1p replacement on all ports:

```
enable dot1p replacement ports all
```

History

This command was available in ExtremeXOS 11.0.

The **qosprofile** keyword and *qosprofile* variable were added in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable eaps

```
enable eaps {name}
```

Description

Enables the EAPS function for a named domain or for an entire switch.

Syntax Description

<i>name</i>	Specifies the name of an EAPS domain.
-------------	---------------------------------------

Default

Disabled.

Default command enables EAPS for the entire switch.

Usage Guidelines



Note

If you use the same name across categories (for example, STPD and EAPS names), you must specify the identifying keyword as well as the actual name.

To configure and enable an EAPS, complete the following steps:

1. Create EAPS domain and assign the name.
2. Configure the control VLAN.
3. Configure the protected VLAN(s).
4. Add the control VLAN to EAPS domain.
5. Add the protected VLAN(s) to EAPS domain.
6. Configure EAPS mode, master or transit.

- 7. Configure EAPS port, secondary and primary.
- 8. If desired, configure timeout and action for failtimer expiration*.
- 9. If desired, configure the hello time for the health-check packets*.
- 10. Enable EAPS for the entire switch.
- 11. If desired, enable Fast Convergence*.
- 12. Enable EAPS for the specified domain.

Although you can enable EAPS prior to configuring these steps, the EAPS domain(s) does not run until you configure these parameters.

* These steps can be configured at any time, even after the EAPS domains are running.

You must enable EAPS globally and specifically for each named EAPS domain.

Example

The following command enables the EAPS function for entire switch:

```
enable eaps
```

The following command enables the EAPS function for the domain eaps-1:

```
enable eaps eaps-1
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable edp ports

```
enable edp ports [ports | all]
```

Description

Enables the EDP on one or more ports.

Syntax Description

<i>ports</i>	Specifies one or more ports or slots and ports, including management port.
all	Specifies all ports on the switch, including management port.

Default

Enabled.

Usage Guidelines

On a SummitStack, *ports* can be a list of slots and ports. On a stand-alone switch, *ports* can be one or more port numbers. For a detailed explanation of port specification, see [Port Numbering](#) in [Command Reference Overview](#)

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 information.
- Switch port number.
- Switch port configuration data: duplex, and speed.

Example

The following command enables EDP on port 3 on a switch:

```
enable edp ports 3
```

History

This command was first available in ExtremeXOS 10.1.

The port configuration data was added in ExtremeXOS 11.0.

Ability to enable EDP on management port was added in ExtremeXOS 30.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable elrp-client

```
enable elrp-client {software | hardware-assist}
```

Description

Enables the Extreme Loop Recovery Protocol (ELRP) client (standalone ELRP) globally.

Syntax Description

software	Select software ELRP (Default).
hardware-assist	Select hardware-assisted ELRP. Not available on the ExtremeSwitching X435 series switches.

Default

By default, ELRP is disabled.

By default, ELRP, when enabled, is software ELRP.

Usage Guidelines

Configure loopback port before enabling hardware-assisted ELRP. Note that hardware-assisted ELRP is not available on the ExtremeSwitching X435 series switches.

The ELRP client must be enabled globally in order for it to work on any [VLANs](#).

The ELRP client can be disabled globally so that none of the ELRP VLAN configurations take effect. Use the `disable elrp-client` command to globally disable the ELRP client.

The ExtremeXOS does not support ELRP and Network Login on the same port.

Example

The following command globally enables the ELRP client:

```
# enable elrp-client
```

The following example enables hardware-assisted ELRP client:

```
# enable elrp-client hardware-assist
```

History

This command was first available in ExtremeXOS 11.1.

Hardware-assisted ELRP option added in ExtremeXOS 30.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable elsm ports

```
enable elsm ports port_list
```

Description

Enables the ELSM protocol for the specified ports.

Syntax Description

<code>port_list</code>	Specifies the port or ports for which ELSM should be enabled.
------------------------	---

Default

The default is disabled.

Usage Guidelines

The ELSM protocol allows you to detect CPU and remote link failures in the network. ELSM operates on a point-to-point basis; you only configure ELSM on the ports that connect to other devices within the network, but you must configure ELSM on both sides of the peer connections.

The Layer 2 connection between the ports determines the peer. You can have a direct connection between the peers or hubs that separate peer ports. In the first instance, the peers are also considered neighbors. In the second instance, the peer is not considered a neighbor.

An Extreme Networks device with ELSM enabled detects CPU and remote link failures by exchanging hello messages between two ELSM peers. If ELSM detects a failure, the ELSM-enabled port responds by blocking traffic on that port. For example, if a peer stops receiving messages from its peer, ELSM brings down that connection by blocking all incoming and outgoing data traffic on the port and notifying applications that the link is down.

When you enable ELSM on a port, ELSM immediately blocks the port and it enters the Down state. When the port detects an ELSM-enabled peer, the peer ports exchange ELSM hello messages. At this point, the ports enter the transitional Down-Wait state. If the port receives Hello+ messages from its peer and does not detect a problem, the peers enter the Up state. If a peer detects a problem or there is no peer port configured, the port enters the Down state.

For more information about the types of ELSM hello messages, see the [configure elsm ports hellotime](#) command.



Note

ELSM and mirroring are mutually exclusive. You can enable either ELSM, or mirroring, but not both.

If you try to enable ELSM on a port that is already configured as a mirrored port, the switch displays a message similar to the following: `Cannot enable ELSM on port 1. Port is configured as mirror monitor port`

Configuring the Hello Timer Interval

ELSM ports use hello messages to communicate information about the health of the network to its peer port. You can also configure the interval by which the ELSM-enabled ports sends hello messages. For more information about configuring the hello interval, see the command [configure elsm ports hellotime](#).

Disabling ELSM

ELSM works between two connected ports, and each ELSM instance is based on a single port. When you disable ELSM on the specified ports, the ports no longer send ELSM hello messages to their peers and no longer maintain ELSM states. To disable ELSM, use the following command:

```
disable elsm ports port_list
```

Example

The following command enables ELSM for ports 1-2 on the switch:

```
enable elsm ports 1-2
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable elsm ports auto-restart

```
enable elsm ports port_list auto-restart
```

Description

Enables ELSM automatic restart for the specified ports.

Syntax Description

<i>port_list</i>	Specifies the port or ports for which ELSM auto-restart is being enabled.
------------------	---

Default

The default is enabled.

Usage Guidelines

You must explicitly configure this behavior on each ELSM-enabled port; this is not a global command.

By default, ELSM automatic restart is enabled. If an ELSM-enabled port goes down, ELSM bypasses the Down-Stuck state and automatically transitions the down port to the Down state, regardless of the number of times the port goes up and down.

If you disable ELSM automatic restart, the ELSM-enabled port can transition between the following states multiple times: Up, Down, and Down-Wait. When the number of state transitions is greater than or equal to the sticky threshold, the port enters the Down-Stuck state.

The ELSM sticky threshold specifies the number of times a port can transition between the Up and Down states. The sticky threshold is not user-configurable and has a default value of 1. That means a port can transition only one time from the Up state to the Down state. If the port attempts a subsequent transition from the Up state to the Down state, the port enters the Down-Stuck state.

If the port enters the Down-Stuck state, you can clear the stuck state and enter the Down state by using one of the following commands:

```
clear elsm ports port_list auto-restart
```

OR

```
enable elsm ports port_list auto-restart
```

If you use the `enable elsm ports port_list auto-restart` command, automatic restart is always enabled; you do not have to use the `clear elsm ports port_list auto-restart` command to clear the stuck state.

To disable automatic restart, use the following command:

```
disable elsm ports port_list auto-restart
```

If you configure automatic restart on one port, Extreme Networks recommends that you use the same configuration on its peer port.

Example

The following command enables ELSM automatic restart for slot 2, ports 1-2 on the switch:

```
enable elsm ports 2:1-2:2 auto-restart
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable erps

```
enable erps
```

Description

Enable (ERPS/ITU-T G.8032 standard).

Syntax Description

N/A.

Default

N/A.

Usage Guidelines

Use this command to enable ERPS.

Example

```
enable erps
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

enable erps block-vc-recovery

```
enable erps ring-name block-vc-recovery
```

Description

Enable ability on ERPS rings to block virtual channel recovery to avoid temporary loops .

Syntax Description

<i>ring-name</i>	Alphanumeric string that identifies the ERPS ring.
block-vc-recovery	Block on Virtual channel recovery.

Default

N/A.

Usage Guidelines

Use this command to enable ability on ERPS rings to block on virtual channel recovery to avoid temporary loops. This is done on interconnected nodes for sub-ring configurations.

Example

The following example enables a virtual channel recovery block on “ring1”:

```
enable erps ring1 block-vc-recovery
```

History

This command was first available in ExtremeXOS 15.13.

Platform Availability

This command is available on all platforms that are running ExtremeXOS.

enable erps ring-name

```
enable erps ring-name
```

Description

Enable an existing *ERPS* ring/sub-ring.

Syntax Description

<i>ring-name</i>	Alphanumeric string that identifies the ERPS ring.
------------------	--

Default

N/A.

Usage Guidelines

Use this command to enable an existing ERPS ring/sub-ring.

Example

The following example enables an existing ERPS ring identified as “ring1”:

```
enable erps ring1
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

enable erps topology-change

enable erps *ring-name* **topology-change**

Description

Enable the ability of *ERPS* to set the topology-change bit to send out Flush events.

Syntax Description

<i>ring-name</i>	Alphanumeric string that identifies the ERPS sub-ring.
topology-change	Topology change propagation control.

Default

N/A.

Usage Guidelines

Use this command to enable the ability of ERPS to set the topology-change bit to send out Flush events.

Example

The following example enables the ability to set the topology-change bit for an existing ERPS sub-ring identified as “ring1”:

```
enable erps ring1 topology-change
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

enable esrp

enable esrp *esrpDomain*

Description

Enables *ESRP* for a named domain.

Syntax Description

<i>esrpDomain</i>	Specifies the name of an ESRP domain.
-------------------	---------------------------------------

Default

Disabled.

Usage Guidelines

Before you enable an ESRP domain, it must have a domain ID. The ESRP domain ID is determined from one of the following user-configured parameters:

- ESRP domain number created with the `configure esrp domain-id` command
- 802.1Q tag (VLANid) of the tagged master *VLAN*

If you do not have a domain ID, you cannot enable ESRP on that domain. A message similar to the following appears:

```
ERROR: Cannot enable ESRP Domain "esrp1" ; No domain id configured!
```

If you add an untagged master VLAN to the ESRP domain, make sure to create an ESRP domain ID with the `configure esrp domain-id` command before you attempt to enable the domain.

Example

The following command enables ESRP for the domain esrp1:

```
enable esrp esrp1
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on all platforms that use the Edge, Advanced Edge, or Core license. For information on the licenses available for each platform, see the [Switch Engine 32.3 Feature License Requirements](#) document.

enable ethernet oam ports link-fault-management

```
enable ethernet oam ports [port_list | all] link-fault-management
```

Description

Enables Ethernet OAM on ports.

Syntax Description

<i>port_list</i>	Specifies the particular ports.
all	Specifies all fiber ports.

Default

Ethernet OAM is disabled on all ports.

Usage Guidelines

Use this command to enable Ethernet OAM on one or more specified ports or on all fiber ports. Unidirectional link fault management is supported only on fiber ports.

Before enabling Ethernet OAM, autonegotiation must be turned off. The link should be a full duplex link.

If some ports cannot be enabled because, for instance, autonegotiation is not turned off, the command is executed for those ports that can be enabled and reasons for the failed ports are displayed.

To display the Ethernet OAM configuration, use the `show ethernet oam` command.

When operating as a stack master, the ExtremeSwitching X450e switch can process this command for ports on supported platforms.

Example

The following command enables Ethernet OAM on all fiber ports:

```
# enable ethernet oam ports all link-fault-management
```

History

This command was first available in ExtremeXOS 12.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable fdb static-mac-move

```
enable fdb static-mac-move
```

Description

Enables EMS and SNMP reporting of discovered MAC addresses that are duplicates of statically configured MAC addresses.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command enables reporting only. All packets that arrive from a duplicate MAC address on another port (other than the statically configured port) are dropped.

The switch reports the source MAC address, port, and VLAN for each duplicate MAC address.

Example

The following command enables this feature:

```
enable fdb static-mac-move
```

History

This command was first available in ExtremeXOS 12.7.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable flooding ports

```
enable flooding [all_cast | broadcast | multicast | unicast] ports  
[port_list | all]
```

Description

Enables egress flooding on one or more ports. You can further identify the type of packets to flood on the specified ports.

Syntax Description

all_cast	Specifies enabling egress flooding for all packets on specified ports.
broadcast	Specifies enabling egress flooding only for broadcast packets.
multicast	Specifies enabling egress flooding only for multicast packets.
unicast	Specifies enabling egress flooding only for unknown unicast packets.
<i>port_list</i>	Specifies one or more ports or slots and ports.
all	Specifies all ports on the switch.

Default

Enabled for all packet types.

Usage Guidelines

Use this command to re-enable egress flooding that you previously disabled using the `disable flooding ports` command.

The following guidelines apply to enabling and disabling egress flooding:

- Disabling multicasting egress flooding does not affect those packets within an *IGMP* membership group at all; those packets are still forwarded out. If IGMP snooping is disabled, multicast packets are not flooded.
- Egress flooding can be disabled on ports that are in a load-sharing group. If that is the situation, the ports in the group take on the egress flooding state of the master port; each member port of the load-sharing group has the same state as the master port.
- *FDB* learning is independent of egress flooding. FDB learning and egress flooding can be enabled or disabled independently.
- Disabling unicast or all egress flooding to a port also stops packets with unknown MAC addresses to be flooded to that port.
- Disabling broadcast or all egress flooding to a port also stops broadcast packets to be flooded to that port.

You can disable egress flooding for unicast, multicast, or broadcast MAC addresses, as well as for all packets on the ports. The default behavior is enabled egress flooding for all packet types.

Example

The following command enables unicast flooding on ports 13-17 on a switch:

```
enable flooding unicast port 13-17
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable flow-control ports

```
enable flow-control [tx-pause {priority priority} | rx-pause {qosprofile qosprofile}] ports [all | port_list]
```

Description

Enables flow control or priority flow control (PFC) on the specified ports.

Syntax Description

tx-pause	Specifies transmit pause frames.
<i>priority</i>	Specifies all priorities or single priorities--dot1p priority for tagged packets and internal priority for untagged packets. Used with priority flow control only.
rx-pause	Specifies received pause frames.
<i>qosprofile</i>	Specifies a QoS profile ("qp1" "qp2" "qp3" "qp4" "qp5" "qp6" "qp7" "qp8") to pause for priority flow control packet reception. Used with priority flow control only.
all	Specifies all ports or slots.
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

Disabled.

Usage Guidelines

With autonegotiation enabled, the switches advertise the ability to support pause frames. This includes receiving, reacting to (stopping transmission), and transmitting pause frames. However, the switch does not actually transmit pause frames unless it is configured to do so.

IEEE 802.3x-Flow Control

IEEE 802.3x flow control provides the ability to configure different modes in the default behaviors.

Use this command to configure the switch to transmit link-layer pause frames when congestion is detected. This stops all traffic on the configured port when there is buffer congestion for any traffic type. Use it also to configure the switch to return to the default behavior of processing received pause frames.

To enable TX flow-control, RX flow-control must first be enabled. If you attempt to enable TX flow-control with RX flow-control disabled, an error message is displayed.

IEEE 802.1Qbb-Priority Flow Control

IEEE 802.1Qbb priority flow control provides the ability to configure the switch to transmit link-layer pause frames to stop only a portion of the traffic when congestion is detected.

When IEEE 802.1Qbb priority flow control is enabled on a port, IEEE 802.3x pause functionality is no longer available on that port.

Priority is established for reception of PFC packets with a QoS profile value on the ExtremeXOS switch and for transmission with a priority value added to the PFC packet.

- QoS profile—Ingress traffic is associated with a QoS profile for assignment to one of eight hardware queues in the system that define how the traffic flows with respect to bandwidth, priority, and other parameters. By default, there are two QoS profiles (QP1 and QP8) defined in these supported

platforms and PFC works with this default. To segregate the ingress traffic with more granularity, you will want to define other QoS profiles.

- Priority—The traffic that is paused is based on the priority bits in the [VLAN](#) header for tagged packets. You can specify this transmit priority independently from the QoS profile to associate it with the reception of a PFC packets thus giving flexibility in the configuration of the network.

It is suggested that the priority in the VLAN header match the QoS profile priority when traffic ingresses at the edge of the network so that the traffic can be more easily controlled as it traverses through the network.

IEEE 802.3x

The following command enables the TX flow-control feature on ports 5 through 7 on a switch:

```
enable flow-control tx-pause ports 5-7
```

IEEE 802.1Qbb

The following command enables the priority flow control feature on a switch:

```
enable flow-control tx-pause priority 3 ports 2
```

History

This command was first available in ExtremeXOS 12.1.3.

IEEE 802.1Qbb priority flow control (PFC) was added in ExtremeXOS 12.5.

Platform Availability

IEEE 802.3x

The basic TX-pause and RX-pause functions of this command are available on all switches.

IEEE 802.1Qbb

The priority function (PFC) is available only on 10G ports.

enable flowmon

```
enable flowmon
```

Description

Enables Flow Monitor to collect and export flows for configured keys in enabled groups.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

Example

The following command enables Flow Monitor:

```
# enable flowmon
```

History

This command was first available in ExtremeXOS 32.2.

Platform Availability

This command is available on ExtremeSwitching 5520 and 5420 series switches.

enable flowmon group

```
enable flowmon group group_name
```

Description

Enables a Flow Monitor group.

Syntax Description

group	Specifies the Flow Monitor group.
<i>group_name</i>	Specifies the assigned name of the Flow Monitor group. Range is 32 characters.

Default

N/A.

Usage Guidelines

While a group is being enabled, it can't be modified except to add or delete keys.

Before a group can be enabled, it must have at least one key and one collector added to it. A template key portion is present once a key has been added.

Example

The following command enables a Flow Monitor group with the name 'max-flow-age':

```
# enable flowmon group max-flow-age
```

History

This command was first available in ExtremeXOS 32.2.

Platform Availability

This command is available on ExtremeSwitching 5520 and 5420 series switches.

enable icmp ipv6

```
enable icmp ipv6 [ignore-multicasts | ignore-anycasts]
```

Description

Enables the *ICMP* IPv6 reply to multicast or anycast echo request.

Syntax Description

ignore-multicasts	Specifies to ignore ICMP echo requests destined to an IP multicast address. Default is ignore (disable).
ignore-anycasts	Specifies to ignore ICMP echo requests destined to an IP anycast address. Default is ignore (disable).

Default

Ignore (disable).

Usage Guidelines

Use this command to enable an ignore reply to multicast or anycast echo request.

Example

The following example specifies to ignore ICMP multicast echo requests:

```
enable icmp ipv6 ignore-multicasts
```

History

This command was first available in ExtremeXOS 31.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable icmp redirects ipv6 fast-path

```
enable icmp redirects ipv6 fast-path
```

Description

When enabled, IPv6 packets forwarded by hardware (fast path) may trigger *ICMP* redirects.

Syntax Description

fast-path	IPv6 packets forwarded by hardware may trigger ICMP redirects
------------------	---

Default

Disabled.

Usage Guidelines

Use this command to trigger ICMP redirects when IPv6 packets are forwarded by hardware (fast-path).

Example

The enabled or disabled setting is displayed when using the command:

```
# show ipconfig ipv6
Route Sharing                : Disabled
ICMP Redirect for Fast Path : Enabled
Max Shared Gateways         : Current: 4   Configured: 4

Interface      IPv6 Prefix      Flags
v1             2001::1/24             -EUf---R-
v1             fe80::204:96ff:fe1e:ec00%v1/64 -EUfP--R-
Flags : D - Duplicate address detected on VLAN, T - Tentative address
E - Interface enabled, U - Interface up, f - IPv6 forwarding enabled,
i - Accept received router advertisements enabled,
R - Send redirects enabled, r - Accept redirects enabled
P - Prefix address
```

History

This command was first available in ExtremeXOS 15.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable icmp redirects

```
enable icmp redirects {ipv4} {vlan all | {vlan} {name}}
```

Description

Enables the generation of *ICMP* redirect messages on one or all *VLANs*.

Syntax Description

<i>name</i>	Specifies a VLAN name.
-------------	------------------------

Default

Enabled.

Usage Guidelines

This option only applies to the switch when the switch is in routing mode.

ICMP redirects are used in the situation where there are multiple routers in the same subnet. If a host sends a packet to one gateway, the gateway router looks at its route table to find the best route to the destination. If it sees that the best route is through a router in the same subnet as the originating host, the switch sends an ICMP redirect (type 5) message to the host that originated the packet, telling it to use the other router with the better route. The switch also forwards the packet to the destination.

ICMP redirects are only generated for IPv4 unicast packets that are "slowpath" forwarded by the CPU. That is, IPv4 packets that contain IP Options, or packets whose Destination IP is not in the Layer 3 forwarding hardware table.

Example

The following example enables the generation of ICMP redirect messages on all VLANs:

```
enable icmp redirects
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable icmp userredirects

```
enable icmp userredirects
```

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

If the switch has a route to a destination network, the switch uses that router as the gateway to forward the packets to. If that router knows about a better route to the destination, and the next hop is in the same subnet as the originating router, the second router sends an ICMP redirect message to the originating router. If ICMP useredirects is enabled, the switch adds a route to the destination network using the third router as the next hop and starts sending the packets to the third router.

Example

The following example enables the modification of route table information:

```
enable icmp useredirects
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable identity-management

```
enable identity-management
```

Description

Enables the identity management feature, which tracks users and devices that connect to the switch.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

Only admin-level users can execute this command.

After identity management is enabled, the software creates two dynamic *ACL* rules named `idm_black_list` and `idm_white_list`. These rules are removed if identity management is disabled.



Note

FDB entries are flushed on identity management enabled ports when this command is executed.

Example

The following command enables the identity management feature:

```
enable identity-management
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable cli idle-timeout

```
enable cli idle-timeout
```

Description

Enables a timer that disconnects Telnet, SSH2, and console sessions after a period of inactivity (20 minutes is default).

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Timeout 20 minutes.

Usage Guidelines

You can use this command to ensure that a Telnet, Secure Shell (SSH2), 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 change the period of inactivity that triggers the timeout for a Telnet, SSH2, or console session, use the `configure timezone` command.

To view the status of idle timeouts on the switch, use the `show management` command. The `show management` command displays information about the switch including the enable/disable state for idle timeouts. You can configure the length of the timeout interval.

Example

The following command enables a timer that disconnects any Telnet, SSH2, and console sessions after 20 minutes of inactivity:

```
enable cli idle-timeout
```

History

This command was first available in ExtremeXOS 10.1.

The `cli` keyword was added and the `idletimeout` keyword was changed to `idle-timeout` in ExtremeXOS 30.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable igmp

```
enable igmp {vlan vlan name } {IGMPv1 | IGMPv2 | IGMPv3}
```

Description

Enables IGMP on a router interface. If no VLAN is specified, IGMP is enabled on all router interfaces.

Syntax Description

<i>vlan name</i>	Specifies a VLAN name.
IGMPv1	Specifies the compatibility mode as IGMPv1.
IGMPv2	Specifies the compatibility mode as IGMPv2.
IGMPv3	Specifies the compatibility mode as IGMPv3.

Default

Enabled, set to IGMPv2 compatibility mode.

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, IP hosts respond to the query, and group registration is maintained.

IGMPv2 is enabled by default on the switch. However, the switch can be configured to disable the generation and processing of IGMP packets. IGMP should be enabled when the switch is configured to perform IP multicast routing.

Example

The following example enables IGMPv2 on the VLAN accounting:

```
enable igmp vlan accounting
```

The following example enables IGMPv3 on the VLAN finance:

```
enable igmp vlan finance igmpv3
```

History

This command was first available in ExtremeXOS 10.1.

The **IGMPv1**, **IGMPv2**, and **IGMPv3** options were added in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IPv4 multicast feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

enable igmp snooping

```
enable igmp snooping {forward-mcrouter-only | {vlan} name | with-proxy  
                    vr vrname}
```

Description

Enables *IGMP* snooping on one or all *VLANs*.

Syntax Description

forward-mcrouter-only	Specifies that the switch forward all multicast traffic to the multicast router only.
<i>name</i>	Specifies a VLAN or VMAN on which to enable IGMP snooping.
with-proxy <i>vr vrname</i>	Controls how join and leave messages are forwarded from the specified virtual router. If this option is specified, one join message per query is forwarded, and a leave message is forwarded only if it is from the last receiver on the VLAN.

Default

Enabled.

Usage Guidelines

This command applies to both IGMPv2 and IGMPv3.

IGMP snooping is enabled by default on the switch. If you are using multicast routing, IGMP snooping can be enabled or disabled. If IGMP snooping is disabled, all IGMP and IP multicast traffic floods within a given VLAN or VMAN.

The **forward-mcrouter-only**, **vlan**, and **with-proxy** options control three separate and independent features. You can manage one feature at a time with this command, and you can enter the command multiple times as needed to control each feature. For example, you can enter the command twice to enable both the **forward-mcrouter-only** and **with-proxy** options.

If a VLAN or VMAN name is specified with this command, IGMP snooping is enabled only on that VLAN or VMAN. If no options are specified, IGMP snooping is enabled on all VLANs.



Note

IGMP snooping is not supported on SVLANs on any platform.

The with-proxy option enables the IGMP snooping proxy feature, which reduces the number of join and leave messages forwarded on the virtual router as described in the table above. This feature is enabled by default.

An optional optimization for IGMP snooping is the strict recognition of routers only if the remote devices are running a multicast protocol. Two IGMP snooping modes are supported:

- The **forward-mcrouter-only** mode forwards all multicast traffic to the multicast router (that is, the router running PIM, DVMRP or CBT).

- When not in the forward-mcrouter-only mode, the switch forwards all multicast traffic to any IP router (multicast or not), and any active member port to the local network that has one or more subscribers.

**Note**

The forward-mcrouter-only mode for IGMP snooping is enabled/disabled on a switch-wide basis, not on a per-VLAN basis. In other words, all the interfaces enabled for IGMP snooping are either in the forward-mcrouter-only mode or in the non-forward-mcrouter-only mode, and not a mixture of the two modes.

To change the snooping mode you must disable IP multicast forwarding. To disable IP multicast forwarding, use the command:

```
disable ipmcforwarding {vlan name}
```

To change the IGMP snooping mode from the non-forward-mcrouter-only mode to the forward-mcrouter-only mode, use the commands:

```
disable ipmcforwarding  
  
enable igmp snooping forward-mcrouter-only  
  
enable ipmcforwarding {vlan name}
```

To change the IGMP snooping mode from the forward-mcrouter-only mode to the non-forward-mcrouter-only mode, use the commands:

```
disable ipmcforwarding  
  
disable igmp snooping forward-mcrouter-only  
  
enable ipmcforwarding {vlan name}
```

Example

The following command enables IGMP snooping on the switch:

```
enable igmp snooping
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IGMP snooping feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

enable igmp snooping vlan fast-leave

```
enable igmp snooping {vlan} name fast-leave
```

Description

Enables the [IGMP](#) snooping fast leave feature on the specified [VLAN](#).

Syntax Description

<i>name</i>	Specifies a VLAN.
-------------	-------------------

Default

Disabled.

Usage Guidelines

The fast leave feature operates only with IGMPv2.

To view the fast leave feature configuration, use the `show configuration msmgr` command. This show command displays the fast leave configuration only when the feature is enabled.

Example

The following example enables the IGMP snooping fast leave feature on the default VLAN:

```
enable igmp snooping "Default" fast-leave
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IGMP snooping feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.
topic/ph ">

enable igmp snooping with-proxy

```
enable igmp snooping with-proxy {{vr} vr_name}
```

Description

Enables the [IGMP](#) snooping proxy. The default setting is enabled.

Syntax Description

<i>vr_name</i>	Specifies a VR.
----------------	-----------------

Default

Enabled.

Usage Guidelines

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 last member leaves the group.

This command can be used for troubleshooting purpose. It should be enabled for normal network operation. The command does not alter the snooping setting.

This feature can be enabled when IGMPv3 is enabled; however, it is not effective for IGMPv3.

Example

The following command enables the IGMP snooping proxy:

```
enable igmp snooping with-proxy
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IGMP snooping feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

enable igmp ssm-map

```
enable igmp ssm-map {vr vr-name}
```

Description

Enables *IGMP* SSM mapping on a VR.

Syntax Description

<i>vr-name</i>	Specifies a virtual router name. If the VR name is omitted, the switch uses the VR specified by the current CLI VR context.
----------------	---

Default

Disabled on all interfaces.

Usage Guidelines

Configure the range of multicast addresses for PIM SSM before you enable IGMP SSM mapping. IGMP SSM mapping operates only with IPv4.

Example

The following example enables IGMP-SSM mapping on the VR in the current CLI VR context:

```
enable igmp ssm-map
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IPv4 multicast feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

enable inline-power

```
enable inline-power [{fast {ports [port_list | all]} } | perpetual]
```

Description

Enables PoE, and perpetual PoE to all ports; or fast PoE to all ports, or selected ports, for some platforms.

Syntax Description

fast	Deliver PoE power to devices at the time of switch power on without waiting for boot up based on last saved PoE state. The default is disabled.
ports	For fast PoE, specifies selecting ports. ExtremeSwitching5520 and 5720 series switches only.
<i>port_list</i>	For fast PoE, specifies the port list separated by a comma or -. ExtremeSwitching 5520 and 5720 series switches only.
all	For fast PoE, specifies selecting all ports. ExtremeSwitching 5520 and 5720 series switches only.
perpetual	Preserves PoE power delivery to devices during reboot. Perpetual PoE is a switch-wide setting. The default is disabled.

Default

By default:

- PoE is enabled.
- Fast PoE is disabled.
- Perpetual PoE is disabled.

Usage Guidelines

You can control whether inline power is provided to the system by using the `disable inline-power` command and the `enable inline-power` command. By default, inline power provided to all ports is enabled. Additionally, you can opt to deliver PoE power to devices at the time of switch power on without waiting for boot up (fast PoE) based on last saved PoE state. Per-port fast PoE is available on certain platforms. You can also elect to preserve PoE power delivery to devices during reboot (perpetual PoE). The default for both PoE options is disabled.

Enabling inline power starts the PoE detection process used to discover, classify, and power remote PDs.

Disabling inline power using the `disable inline-power` command does not affect the data traffic traversing the port. And, disabling the port using the `disable port` command does not affect the inline power supplied to the port.



Note

Inline power cannot be delivered to connected PDs unless the switch is powered on.

Example

The following command enables inline power currently provided to all ports:

```
# enable inline-power
```

The following example turns on perpetual PoE for the switch:

```
# enable inline-power perpetual
```

The following example turns on fast PoE for ports 1,2, and 5:

```
# enable inline-power fast ports 1,2,5
```

History

This command was first available in ExtremeXOS 11.1.

The **fast** and **perpetual** PoE options were added in ExtremeXOS 30.3.

Per-port fast PoE was added for ExtremeXOS 31.1.

Platform Availability

This command is available on the PoE devices listed in [Extreme Networks PoE Devices](#).

The **fast** and **perpetual** options are only available on the ExtremeSwitching 5520 and 5720 (per port) series switches.

enable inline-power ports

```
enable inline-power ports [all | port_list]
```

Description

Enables *PoE* power currently provided to all ports or to specified ports.

Syntax Description

all	Enables inline power to all ports on the switch.
<i>port_list</i>	Enables inline power to the specified ports.

Default

Enable.

Usage Guidelines

Disabling inline power to a port immediately removes power to any connected PD. By default, inline power provided to all ports is enabled.

Disabling inline power using the `disable inline-power` command does not affect the data traffic traversing the port. And, disabling the port using the `disable port` command does not affect the inline power supplied to the port.

Example

The following command enables inline power to ports 4 and 5 on a switch:

```
enable inline-power ports 4-5
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on the PoE devices listed in [Extreme Networks PoE Devices](#).

enable inline-power slot

```
enable inline-power [fast | perpetual] slot slot
```


Description

Enables *PoE* power, and fast and perpetual PoE power to the specified node (slot) on SummitStacks.

Syntax Description

<i>slot</i>	Enables inline power to specified slot.
fast	Deliver PoE power to devices at the time of switch power on without waiting for boot up based on last saved PoE state. The default is disabled.
perpetual	Preserves PoE power delivery to devices during reboot. Perpetual PoE is a switch-wide setting. The default is disabled.

Default

By default:

- PoE is enabled.
- Fast PoE is disabled.
- Perpetual PoE is disabled.

Usage Guidelines

Disabling inline power to a slot immediately removes power to any connected PDs. By default, inline power provided to all slots is enabled. Additionally, you can opt to deliver PoE power to devices at the time of switch power on without waiting for boot up (fast PoE) based on last saved PoE state. You can also elect to preserve PoE power delivery to devices during reboot (perpetual PoE). The default for both PoE options is disabled.

To deliver inline power to slots, you must reserve power for that slot using the `configure inline-power budget` command. By default, each PoE module has 50 W of power reserved for inline power.

Disabling inline power using the `disable inline-power` command does not affect the data traffic traversing the slot. And, disabling the slot using the `disable slot` command does not affect the inline power supplied to the slot.

If you do not specify a slot number, the command operates on all active nodes in the stack. This command operates only on nodes in the active topology.

Example

The following command makes inline power available to slot 3:

```
# enable inline-power slot 3
```

The following example turns on perpetual PoE to slot 3:

```
# enable inline-power perpetual slot 3
```

History

This command was first available in ExtremeXOS 11.1.

The **fast** and **perpetual** PoE options were added in ExtremeXOS 30.3.

Platform Availability

This command is available on SummitStack when the stack contains switches listed in [Extreme Networks PoE Devices](#).

enable ip anycast

```
enable ip anycast {vlan} vlan_name
```

Description

Enables IP anycast on a VLAN.

Syntax Description

ip	Layer 3 Internet Protocol.
anycast	Enables IP anycast-enabled VLANs to use the anycast MAC on that VLAN.
vlan	Selects the VLAN.
<i>vlan_name</i>	Specifies the VLAN name.

Default

N/A.

Usage Guidelines

To configure the anycast gateway MAC address that is used by VLANs that enable IP anycast, use the command configure **ip anycast mac** [**none** | *mac*].

Example

The following example enables IP anycast on the VLAN "vlan1":

```
# enable ip anycast vlan vlan1
```

History

This command was first available in ExtremeXOS 30.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable ip nat

```
enable ip nat
```

Description

Globally enables Network Address Translation (NAT).

Syntax Description

ip	Specifies Internet Protocol (IP).
nat	Specifies enabling NAT.

Default

N/A.

Usage Guidelines

NAT maps IP addresses from one address domain (typically private IP address space) to an another address domain (typically a public Internet IP address space) to provide transparent routing to end hosts. This translation is accomplished transparently by having a NAT device translate the IP address and/or Layer 4 port of the packets.

To view IP NAT information, run the command `show ip nat`.

Example

The following example enables IP NAT:

```
# enable ip nat
```

History

This command was first available in ExtremeXOS 31.2.

Platform Availability

This command is available on the switches that support the IP NAT feature. For information about which switches support this and other features, see the [Switch Engine 32.3 Feature License Requirements](#) document.

enable ip nat rule

```
enable ip nat rule rule_name
```

Description

Enables Network Address Translation (NAT) rules.

Syntax Description

ip	Specifies Internet Protocol (IP).
nat	Specifies NAT.
rule	Specifies enabling a NAT rule.
<i>rule_name</i>	Specifies the NAT rule to enable.

Default

N/A.

Usage Guidelines

A rule is programmed in hardware only after global NAT and the specific rule are enabled.

Enabling a rule is allowed only after all of the necessary configurations for the rule are complete. After a rule is enabled, configuration changes (IP address, egress VLAN, etc.) are not allowed for the rule. You must disable a rule to make configuration changes to it.

Example

The following example enables the IP NAT rule "rule1":

```
# enable ip nat rule rule1
```

History

This command was first available in ExtremeXOS 31.2.

Platform Availability

This command is available on the switches that support the IP NAT feature. For information about which switches support this and other features, see the [Switch Engine 32.3 Feature License Requirements](#) document.

enable iparp checking

```
enable iparp {vr vr_name} checking
```

Description

Enables checking if the ARP request source IP address is within the range of the local interface or VLAN domain.

Syntax Description

<i>vr_name</i>	Specifies a VR or VRF.
----------------	------------------------

Default

Enabled.

Usage Guidelines

If you do not specify a VR or VRF, the command applies to the current VR context.

Example

The following example enables IP ARP checking:

```
enable iparp checking
```

History

This command was first available in ExtremeXOS 10.1.

The **vr** option was added in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable iparp gratuitous protect

```
enable iparp gratuitous protect [ {vlan} vlan_name | vlan vlan_list]
```

Description

Enables gratuitous ARP protection on the specified VLAN.

Syntax Description

<i>vlan_name</i>	Specifies the VLAN.
<i>vlan_list</i>	Specifies a VLAN list of IDs.

Default

By default, gratuitous ARP is disabled.

Usage Guidelines

Beginning with ExtremeXOS 11.6, this command replaces this command for configuring gratuitous ARP.

Hosts can launch man-in-the-middle attacks by sending out gratuitous ARP requests for the router's IP address. This results in hosts sending their router traffic to the attacker, and the attacker forwarding that data to the router. This allows passwords, keys, and other information to be intercepted.

To protect against this type of attack, the router will send out its own gratuitous ARP request to override the attacker whenever a gratuitous ARP broadcast with the router's IP address as the source is received on the network.

Example

The following example enables gratuitous ARP protection for VLAN corp:

```
enable iparp gratuitous protect vlan corp
```

History

This command was first available in ExtremeXOS 11.2.

The *vlan_list* option was added in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable iparp refresh

```
enable iparp {vr vr_name} refresh
```

Description

Enables IP ARP to refresh its IP ARP entries before timing out.

Syntax Description

<i>vr_name</i>	Specifies a VR or VRF.
----------------	------------------------

Default

Enabled.

Usage Guidelines

If ARP refresh is enabled, the switch resends ARP requests for the host at 3/4 of the configured ARP timer value.

For example: If the ARP timeout is set to 20 minutes, the switch attempts to resend an ARP request for the host when the host entry is at 15 minutes. If the host replies, the ARP entry is reset back to 0, and the timer starts again.

If you do not specify a VR or VRF, the command applies to the current VR context.

Example

The following example enables IP ARP refresh:

```
enable iparp refresh
```

History

This command was first available in ExtremeXOS 10.1.

The **vr** option was added in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable ipforwarding ipv6

```
enable ipforwarding ipv6 [ {vlan} vlan_name | vlan vlan_list] | tunnel  
    tunnel_name | vr vr_name}
```

Description

Enables IPv6 routing VLANs. If no argument is provided, enables IPv6 routing for all VLANs and tunnels that have been configured with an IPv6 address on the current VR or VRF.

Syntax Description

<i>vlan_name</i>	Specifies an IPv6 configured VLAN.
<i>tunnel_name</i>	Specifies an IPv6 tunnel.
<i>vr_name</i>	Specifies a VR or VRF.

Default

Disabled.

Usage Guidelines

When new IPv6 interfaces are added, IPv6 forwarding is disabled by default.

Example

The following example enables forwarding of IPv6 traffic for all VLANs in the current VR context with IPv6 addresses:

```
enable ipforwarding ipv6
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

enable ipforwarding

```
enable ipforwarding {ipv4 | broadcast} {vlan vlan_name}
```

Description

Enables IPv4 routing or IPv4 broadcast forwarding for one or all VLANs. If no argument is provided, enables IPv4 routing for all VLANs that have been configured with an IP address on the current VR or VRF.

Syntax Description

ipv4	Specifies IPv4 forwarding.
broadcast	Specifies broadcast IP forwarding.
<i>vlan_name</i>	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.

The broadcast, ignore-broadcast, and fast-directbroadcast options each prompt with a warning message when executed while the IP forwarding on the corresponding VLAN is disabled. The hardware and software are NOT programmed until IP forwarding is enabled on the VLAN.

The fast-direct-broadcast and ignore-broadcast options cannot be enabled simultaneously. These are mutually exclusive.

The broadcast option can be enabled in conjunction with fast-direct-broadcast and ignore-broadcast.

Example

The following command enables forwarding of IP traffic for all VLANs in the current VR context 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 first available in ExtremeXOS 10.1.
- The ipv4 keyword was added in ExtremeXOS 11.2.
- The ignore-broadcast and the fast-direct-broadcast keywords were added in ExtremeXOS 12.0.

Platform Availability

This command is available on all platforms that use the Edge, Advanced Edge, or Core license. For information on the licenses available for each platform, see the [Switch Engine 32.3 Feature License Requirements](#) document.

enable ipmcforwarding ipv6

```
enable ipmcforwarding ipv6 {vlan name }
```

Description

Enables IPv6 multicast forwarding on a router interface.

Syntax Description

<i>name</i>	Specifies a <u>VLAN</u> name.
-------------	-------------------------------

Default

Disabled.

Usage Guidelines

If no options are specified, all configured IPv6 interfaces are affected. When new IPv6 interfaces are created, IPv6 multicast forwarding is disabled by default.

IPv6 forwarding must be enabled before enabling IPv6 multicast forwarding.

Example

The following example enables IPv6 multicast forwarding on VLAN accounting:

```
enable ipmcforwarding ipv6 vlan accounting
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IPv6 multicast feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

enable ipmcforwarding

```
enable ipmcforwarding {vlan name}
```

Description

Enables IP multicast forwarding on an IP interface.

Syntax Description

<i>name</i>	Specifies a <u>VLAN</u> 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.

Example

The following example enables IPMC forwarding on the VLAN accounting:

```
enable ipmcforwarding vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IPv4 multicast feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

NEW! enable ipmcforwarding tunnel

```
enable ipmcforwarding {tunnel tunnel_name}
```

Description

Enables IP multicast tunnel forwarding on an IP interface.

Syntax Description

<i>tunnel_name</i>	Specifies a tunnel 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.

Example

The following example enables IPMC forwarding on the tunnel accounting:

```
enable ipmcforwarding tunnel accounting
```

History

This command was first available in ExtremeXOS 32.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IPv4 multicast feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

enable ip option loose-source-route

```
enable ip option loose-source-route
```

Description

Enables processing of the loose source route IP option in the IPv4 packet header.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

This enables the switch to forward IP packets that have the loose source route IP option (0x83) enabled.

Source routing is used when a sending host specifies the router interfaces that the packet must traverse on its way to its destination.

With loose source routing enabled, the packet is forwarded if the routing table has a reverse path to the source IP address of the packet.

Example

The following command enables processing of the loose source route IP option:

```
# enable ip-option loose-source-route
```

History

This command was first available in ExtremeXOS 10.1.

This command was removed in ExtremeXOS 30.1, and then re-introduced in ExtremeXOS 31.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable ip option strict-source-route

```
enable ip option strict-source-route
```

Description

Enables processing of the strict source route IP option in the IPv4 packet header.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

This enables the switch to forward IP packets that have the strict source route IP option (0x89) enabled.

Source routing is used when a sending host specifies the router interfaces that the packet must traverse on its way to its destination.

When strict source routing is used, it means that the packet must use the exact path of routers that lie in the designated router path.

With strict source routing enabled, the switch forwards IP packets with the strict source route option enabled, only if the switch's IP is in the designated list and as long as the next hop in the list is directly attached to one of the router's interfaces.

Example

The following example enables processing of the strict source route IP option:

```
# enable ip-option strict-source-route
```

History

This command was first available in ExtremeXOS 10.1.

This command was removed in ExtremeXOS 30.1, and then re-introduced in ExtremeXOS 31.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable iproute bfd

```
enable iproute bfd {gateway} ip_addr {vr vrname}
```

Description

Enables the BFD client to provide services for IPv4 static routes.

Syntax Description

<i>ip_addr</i>	Specifies the IPv4 address of a neighbor to which BFD services are to be provided.
<i>vrname</i>	Specifies the VR or VRF name for which BFD services are being enabled.

Default

Disabled.

Usage Guidelines

To enable BFD services to an IPv4 neighbor, you must do the following:

- Execute this command on the switches at both ends of the link.
- Enable BFD for specific IPv4 static routes with the `configure iproute add [destination network] [gateway] bfd` command.

Once a BFD session is established between two neighbors, BFD notifies the Route Manager process of the BFD session status and any changes. If other BFD clients (such as the *MPLS* BFD client) are configured between the same neighbors, the clients share a single session between the neighbors.

Example

The following example enables BFD client protection for communications with neighbor 10.10.10.1:

```
# enable iproute bfd 10.10.10.1
```

History

This command was first available in ExtremeXOS 12.5.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable iproute bfd strict

```
enable iproute {protection} bfd strict
```

Description

Turns on "strict" Bidirectional Forwarding Detection (BFD) session control, which brings down the static route during switch reboot if the static route nexthop BFD session is in the INIT state.

Syntax Description

protection	Enables or disables route protection.
bfd	BFD protect static routes to next hop gateway.
strict	Enables considering that protected static routes are not up if the BFD session is in INIT state. Default is disabled.

Default

By default, strict BFD session control is disabled.

Usage Guidelines

If the BFD session is down, but BFD protected static route is still in the routing table after reboot, the BFD session is never established, because during reboot, the BFD session is in the INIT state, and the static route is brought up without considering BFD session state. This can cause traffic loss since the link to the gateway actually is down. This command turns down the static route during reboot if BFD session is in the INIT state. This behavior is different from other BFD clients (such as *OSPF*) in the same INIT situation. A reboot is required to make the command take effect.

Example

The following example enables BFD strict session control:

```
# enable iproute bfd strict
WARNING: Please reboot the switch for the strict BFD to take effect.
```

History

This command was first available in ExtremeXOS 30.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable iproute compression

```
enable iproute compression {vr vrname}
```

Description

Enables IPv4 route compression.

Syntax Description

<i>vrname</i>	VR or VRF name for which the IP route compression is being enabled.
---------------	---

Default

Enabled.

Usage Guidelines

Enables IPv4 route compression for the specified VR or VRF. If the VR name is not specified, route compression is enabled for the VR context from which the CLI command is issued.

The command applies a compression algorithm on each of the IP prefixes in the routing table. Essentially, routes with longer network masks might not be necessary if they are a subset of other routes with shorter network masks using the same gateway(s). When IP route compression is enabled, these unnecessary routes are not provided to the Forwarding Information Base (FIB).

Example

The following example enables IP route compression:

```
enable iproute compression
```

History

This command was first available in ExtremeXOS 12.0.

Default changed to enabled in ExtremeXOS 15.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable iproute ipv6 compression

```
enable iproute ipv6 compression {vr vrname}
```

Description

Enables IPv6 route compression.

Syntax Description

<i>vrname</i>	Specifies a VR or VRF.
---------------	------------------------

Default

If no VR name is specified, the current CLI context is used.

Usage Guidelines

This command enables IPv6 route compression for the VR. This command applies a compression algorithm to each IPv6 prefix in the IPv6 prefix database.

Example

The following example enables IPv6 route compression in the current VR context.

```
enable iproute ipv6 compression
```

History

This command was first available in ExtremeXOS 12.0.

Default changed to enabled in ExtremeXOS 15.6.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

enable iproute mpls-next-hop

```
enable iproute mpls-next-hop
```

Description

Enables IP forwarding over MPLS LSPs for the default VR.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command enables IP forwarding over MPLS LSPs for the default VR. When enabled, LSP next hops can be used to tunnel IP traffic across the MPLS network. By default, IP forwarding over MPLS LSPs is disabled.

Example

The following command enables IP forwarding over MPLS LSPs:

```
enable iproute mpls-next-hop
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

enable iproute protection ping

```
enable iproute {ipv4 | ipv6} protection ping
```

Description

Globally enables ping protection for static routes added with ping protection for IPv4 and IPv6.

Syntax Description

ipv4	Specifies IPv4 (default).
ipv6	Specifies IPv6.
protection	Enables route protection.
ping	Globally enables ping protection for static routes added with ping protection (default is enabled).

Default

Enabled is the default. If not specified, IPv4 is the default.

Example

The following example enables ping protection for static routes added with ping protection for IPv4:

```
# enable iproute ipv4 protection ping
```

History

This command was first available in ExtremeXOS 22.1.

Platform Availability

This command is available on all platforms with any license level as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

enable iproute sharing

```
enable iproute {ipv4 | ipv6} sharing {{vr} vrname} | {{vr} all}}
```

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 are shared.

Syntax Description

<code>vrname</code>	VR or VRF name for which IP route sharing is being enabled.
---------------------	---

Default

Disabled.

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](#), [BGP](#), or IS-IS routes. In OSPF, BGP, and IS-IS, this capability is referred to as [ECMP](#) routing.

Configure static routes and OSPF, BGP, or IS-IS as you would normally. The ExtremeXOS software supports route sharing across up to 64 way ECMP for OSPFv2, BGP, and static routes, or up to 64-way ECMP for [OSPFv3](#) and 8 way for ISIS. However, on SummitStack, and ExtremeSwitching series switches, by default, up to four routes are supported. To support 2, 4, 8, 16, 32, or 64 routes on these switches, use the following command:

```
configure iproute sharing max-gateways max_gateways
```

If a VR is not specified, this command enables IP route sharing in the current VR context.

Example

The following example enables load sharing for multiple routes:

```
enable iproute sharing
```

History

This command was first available in ExtremeXOS 11.1.

The **vr** option was added in ExtremeXOS 12.5.

The **ipv6** option was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable ip-security anomaly-protection icmp

```
enable ip-security anomaly-protection icmp {slot [ slot | all ] }
```

Description

Enables *ICMP* size and fragment checking.

Syntax Description

<i>slot</i>	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is disabled.

Usage Guidelines

This command enables ICMP size and fragment checking. This checking takes effect for both IPv4 and IPv6 TCP packets. When enabled, the switch drops ICMP packets if one of following condition is true:

- Fragmented ICMP packets.
- IPv4 ICMP pings packets with payload size greater than the maximum IPv4 ICMP-allowed size. (The maximum allowed size is configurable.)
- IPv6 ICMP ping packets with payload size > the maximum IPv6 ICMP-allowed size. (The maximum allowed size is configurable.)

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable ip-security anomaly-protection ip

```
enable ip-security anomaly-protection ip { slot [ slot | all ] }
```

Description

Enables source and destination IP address checking.

Syntax Description

<i>slot</i>	Specifies the slot.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is disabled.

Usage Guidelines

This command enables source and destination IP addresses checking. This checking takes effect for both IPv4 and IPv6 packets. When enabled, the switch drops IPv4/IPv6 packets if its source IP address are the same as the destination IP address. In most cases, the condition of source IP address being the same as the destination IP address indicates a Layer 3 protocol error. (These kind of errors are found in LAND attacks.)

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable ip-security anomaly-protection l4port

```
enable ip-security anomaly-protection l4port [tcp | udp | both] {slot  
[ slot | all ]}
```

Description

Enables TCP and UDP ports checking.

Syntax Description

tcp	Specifies that the TCP port be enabled for checking.
udp	Specifies that the UDP port be enabled for checking.
both	Specifies both the TCP and UDP ports be enabled for checking.
<i>slot</i>	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is disabled.

Usage Guidelines

This command enabled TCP and UDP ports checking. This checking takes effect for both IPv4 and IPv6 TCP and UDP packets. When enabled, the switch drops TCP and UDP packets if its source port is the same as its destination port. In most cases, when the condition of source port is the same as that of the destination port, it indicates a Layer4 protocol error. (This type of error can be found in a BALT attack.)

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable ip-security anomaly-protection notify

```
enable ip-security anomaly-protection notify [log | snmp | cache] {slot
[ slot | all ] }
```

Description

Enables protocol anomaly notification.

Syntax Description

log	Specifies the switch to send the notification to a log file.
snmp	Specifies the switch to send an SNMP trap when an event occurs.
cache	Specifies the switch to send the notification to cache.
<i>slot</i>	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is disabled.

Usage Guidelines

This command enables anomaly notification. When enabled, any packet failed to pass enabled protocol checking is sent to XOS Host CPU and notifies the user. There are three different types of notifications:

- **log**: The anomaly events are logged into EMS log.
- **snmp**: The anomaly events generate SNMP traps.
- **cache**: The most recent and unique anomaly events are stored in memory for review and investigation.

When disabled, the switch drops all violating packets silently.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable ip-security anomaly-protection tcp flags

```
enable ip-security anomaly-protection tcp flags {slot [ slot | all ]}
```

Description

Enables TCP flag checking.

Syntax Description

<i>slot</i>	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is disabled.

Usage Guidelines

This command Enables TCP flag checking. This checking takes effect for both IPv4 and IPv6 TCP packets. When enabled, the switch drops TCP packets if one of following condition is true:

- TCP SYN flag==1 and the source port<1024
- TCP control flag==0 and the sequence number==0
- TCP FIN, URG, and PSH bits are set, and the sequence number==0
- TCP SYN and FIN both are set.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable ip-security anomaly-protection tcp fragment

```
enable ip-security anomaly-protection tcp fragment {slot [ slot | all ]}
```

Description

Enables TCP fragment checking.

Syntax Description

<i>slot</i>	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is disabled.

Usage Guidelines

This command enables TCP fragment checking. This checking takes effect for IPv4/IPv6. When it is enabled, the switch drops TCP packets if one of following condition is true:

- For the first IPv4 TCP fragment (its IP offset field==0), if its TCP header is less than the minimum IPv4 TCP header allowed size.
- For the first IPv6 TCP fragment (its IP offset field==0), if its TCP header is less than the minimum IPv6 TCP header allowed size.
- If its IP offset field==1 (for IPv4 only).

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable ip-security anomaly-protection

```
enable ip-security anomaly-protection {slot [ slot | all ] }
```

Description

Enables all anomaly checking options.

Syntax Description

<i>slot</i>	Specifies the slot to be used.
all	Specifies all IP addresses, or all IP addresses in a particular state.

Default

The default is disabled.

Usage Guidelines

This command enables all anomaly checking options, including IP address, UDP/TCP port, TCP flag and fragment, and *ICMP* anomaly checking.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable ip-security arp gratuitous-protection

```
enable ip-security arp gratuitous-protection [dynamic | {vlan} all |  
vlan_name]
```

Description

Enables gratuitous ARP protection on one or all *VLANs* on the switch.

Syntax Description

all	Specifies all VLANs configured on the switch.
<i>vlan_name</i>	Specifies the VLAN.
dynamic	Configuration options for dynamically created VLANs.

Default

By default, gratuitous ARP protection is disabled.

Dynamic VLAN option added in ExtremeXOS 30.2.

Usage Guidelines

Beginning with ExtremeXOS 11.6, this command replaces the `enable iparp gratuitous protect` command.

Hosts can launch man-in-the-middle attacks by sending out gratuitous ARP requests for the router's IP address. This results in hosts sending their router traffic to the attacker, and the attacker forwarding that data to the router. This allows passwords, keys, and other information to be intercepted.

To protect against this type of attack, the router will send out its own gratuitous ARP request to override the attacker whenever a gratuitous ARP broadcast with the router's IP address as the source is received on the network.

Beginning with ExtremeXOS 11.6, if you enable both DHCP secured ARP and gratuitous ARP protection, the switch protects its own IP address and those of the hosts that appear as secure entries in the ARP table.

To protect the IP addresses of the hosts that appear as secure entries in the ARP table, use the following commands to enable DHCP snooping, DHCP secured ARP, and gratuitous ARP on the switch:

- `enable ip-security dhcp-snooping {vlan} vlan_name ports [all | ports] violation-action [drop-packet {[block-mac | block-port] [durationduration_in_seconds | permanently] | none}] {snmp-trap}`
- `enable ipsecurity arp learning learn-from-arp`
- `enable ip-security arp gratuitous-protection {vlan} [all | vlan_name]`

Displaying Gratuitous ARP Information

To display information about gratuitous ARP, use the following command:

```
show ip-security arp gratuitous-protection
```

Example

The following command enables gratuitous ARP protection for VLAN corp:

```
enable ip-security arp gratuitous-protection vlan corp
```

History

This command was first available in ExtremeXOS 11.6.

Dynamic VLAN option added in ExtremeXOS 30.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable ip-security arp learning learn-from-arp

```
enable ip-security arp learning learn-from-arp [dynamic | {vlan}  
vlan_name] ports [all | ports]
```

Description

Enables ARP learning for the specified VLAN and member ports.

Syntax Description

<i>vlan_name</i>	Specifies the name of the VLAN to which this rule applies.
dynamic	Configuration options for dynamically created VLANs.
all	Specifies all ingress ports.
<i>ports</i>	Specifies one or more ingress ports.

Default

By default, ARP learning is enabled.

Usage Guidelines

ARP is part of the TCP/IP suite used to associate a device's physical address (MAC address) with its logical address (IP address). The switch broadcasts an ARP request that contains the IP address, and the device with that IP address sends back its MAC address so that traffic can be transmitted across the network. The switch maintains an ARP table (also known as an ARP cache) that displays each MAC address and its corresponding IP address.

By default, the switch builds its ARP table by tracking ARP requests and replies, which is known as ARP learning.

Displaying ARP Information

To display how the switch builds an ARP table and learns MAC addresses for devices on a specific VLAN and associated member ports, use the following command:

```
show ip-security arp learning {vlan} vlan_name
```

To view the ARP table, including permanent and *DHCP* secured ARP entries, use the following command:

```
show iparp {ip_addre | mac | vlanvlan_name | permanent} {vrvr_name}
```

Example

The following command enables ARP learning on port 1:1 of the VLAN learn:

```
enable ip-security arp learning learn-from-arp vlan learn ports 1:1
```

History

This command was first available in ExtremeXOS 11.6.

Dynamic VLAN option added in ExtremeXOS 30.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable ip-security arp learning learn-from-dhcp

```
enable ip-security arp learning learn-from-dhcp [dynamic vlan | {vlan}  
vlan_name] ports [all | ports]
```

Description

Enables *DHCP* secured ARP learning for the specified *VLAN* and member ports.

Syntax Description

dynamic	Configuration options for dynamically created VLANs.
<i>vlan_name</i>	Specifies the name of the VLAN to which this rule applies.
all	Specifies all ingress ports.
<i>ports</i>	Specifies one or more ingress ports.

Default

By default, DHCP secured ARP learning is disabled.

Usage Guidelines

Use this command to configure the switch to add the MAC address and its corresponding IP address to the ARP table as a secure ARP entry. The switch does not update secure ARP entries, regardless of the ARP requests and replies seen by the switch. DHCP secured ARP is linked to the “DHCP snooping” feature. The same DHCP bindings database created when you enabled DHCP snooping is also used by DHCP secured ARP to create secure ARP entries. The switch only removes secure ARP entries when the corresponding DHCP entry is removed from the trusted DHCP bindings database.



Note

If you enable DHCP secured ARP on the switch, ARP learning continues, which allows insecure entries to be added to the ARP table.

The default ARP timeout (configure iparp timeout) and ARP refresh (enable iparp refresh) settings do not apply to DHCP secured ARP entries. The switch removes DHCP secured ARP entries upon any DHCP release packet received from the DHCP client.

Displaying ARP Information

To display how the switch builds an ARP table and learns MAC addresses for devices on a specific VLAN and associated member ports, use the following command:

```
show ip-security arp learning {vlan} vlan_name
```

To view the ARP table, including permanent and DHCP secured ARP entries, use the following command:

```
show iparp {ip_address | mac | vlanvlan_name | permanent} {vrvr_name}
```

Example

The following command enables DHCP secured ARP learning on port 1:1 of the VLAN learn and uses the default polling and retry intervals:

```
enable ip-security arp learning learn-from-dhcp vlan learn ports 1:1
```

History

This command was first available in ExtremeXOS 11.6.

Dynamic VLAN support was added in ExtremeXOS 30.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable ip-security arp validation violation-action

```
enable ip-security arp validation {destination-mac} {source-mac} {ip}
  [dynamic vlan_id | {vlan} vlan_name] [all | ports] violation-
action [drop-packet {[block-port] [duration duration_in_seconds |
permanently]}] {snmp-trap}
```

Description

Enables ARP validation for the specified VLAN and member ports.

Syntax Description

destination-mac	Specifies that the switch checks the ARP payload for the MAC destination address in the Ethernet header and the receiver's host address in the ARP response.
source-mac	Specifies that the switch checks ARP requests and responses for the MAC source address in the Ethernet header and the sender's host address in the ARP payload.
ip	Specifies the switch checks the IP address in the ARP payload and compares it to the DHCP bindings database. If the IP address does exist in the DHCP bindings table, the switch verifies that the MAC address is the same as the sender hardware address in the ARP request. If not, the packet is dropped.
dynamic	Configuration options for dynamically created VLANs.
<i>vlan_id</i>	VLAN ID tag between 1 and 4,094.
<i>vlan_name</i>	Specifies the name of the VLAN to which this rule applies.
all	Specifies all ports to participate in ARP validation.
<i>ports</i>	Specifies one or more ports to participate in ARP validation.
drop-packet	Specifies that the switch drops the invalid ARP packet.

block-port	Indicates that the switch blocks invalid ARP requests on the specified port.
<i>duration_in_seconds</i>	Specifies the switch to temporarily disable the specified port upon receiving an invalid ARP request. The range is seconds.
permanently	Specifies the switch to permanently disable the port upon receiving an invalid ARP request.
snmp-trap	Specifies the switch to send an SNMP trap when an event occurs.

Default

By default, ARP validation is disabled.

Usage Guidelines

The violation action setting determines what action(s) the switch takes when an invalid ARP is received.

Depending on your configuration, the switch uses the following methods to check the validity of incoming ARP packets:

- **Drop packet**—The switch confirms that the MAC address and its corresponding IP address are in the DHCP binding database built by DHCP snooping. This is the default behavior when you enable ARP validation. If the MAC address and its corresponding IP address are in the DHCP bindings database, the entry is valid. If the MAC address and its corresponding IP address are not in the DHCP bindings database, the entry is invalid, and the switch drops the ARP packet.
- **IP address**—The switch checks the IP address in the ARP payload. If the switch receives an IP address in the ARP payload that is in the DHCP binding database, the entry is valid. If the switch receives an IP address that is not in the DHCP binding database, for example 255.255.255.255 or an IP multicast address, the entry is invalid or unexpected.
- **Source MAC address**—The switch checks ARP requests and responses for the source MAC address in the Ethernet header and the sender's host address in the ARP payload. If the source MAC address and sender's host address are the same, the entry is valid. If the source MAC source and the sender's host address are different, the entry is invalid.
- **Destination MAC address**—The switch checks the ARP payload for the destination MAC address in the Ethernet header and the receiver's host address. If the destination MAC address and the target's host address are the same, the entry is valid. If the destination MAC address and the target's host address are different, the entry is invalid.

Any violation that occurs causes the switch to generate an EMS log message. You can configure to suppress the log messages by configuring EMS log filters.

Displaying ARP Validation Information

To display information about ARP validation, use the following command:

```
show ip-security arp validation {vlan} vlan_name
```

Example

The following example enables ARP validation on port 1:1 of the VLAN valid:

```
enable ip-security arp validation vlan valid ports 1:1 drop-packet
```

History

This command was first available in ExtremeXOS 11.6.

Dynamic VLAN and VLAN ID options added in ExtremeXOS 30.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable ip-security dhcp-bindings restoration

```
enable ip-security dhcp-bindings restoration
```

Description

Enables download and upload of *DHCP* bindings.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

The command allows you to enable the download and upload of the DHCP bindings, essentially enabling the DHCP binding functionality. The default is disabled.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable ip-security dhcp-snooping

```
enable ip-security dhcp-snooping [dynamic | {vlan} vlan_name] ports
    [all | ports] violation-action [drop-packet {[block-mac | block-port]
    [duration duration_in_seconds | permanently] | none]}] {snmp-trap}
```

Description

Enables DHCP snooping for the specified VLAN and ports.

Syntax Description

<i>vlan_name</i>	Specifies the name of the DHCP-snooping VLAN. Create and configure the VLAN before enabling DHCP snooping.
dynamic	Configuration options for dynamically created VLANs.
all	Specifies all ports to receive DHCP packets.
<i>ports</i>	Specifies one or more ports to receive DHCP packets.
drop-packet	Indicates that the switch drop the rogue DHCP packet received on the specified port.
block-mac	Indicates that the switch blocks rogue DHCP packets from the specified MAC address on the specified port. The MAC address is added to the DHCP bindings database.
block-port	Indicates that the switch blocks rogue DHCP packets on the specified port. The port is added to the DHCP bindings database.
<i>duration_in_seconds</i>	Specifies that the switch temporarily disable the specified port upon receiving a rogue DHCP packet. The range is seconds.
permanently	Specifies that the switch to permanently disable the specified port upon receiving a rogue DHCP packet.
none	Specifies that the switch takes no action when receiving a rogue DHCP packet; the switch does not drop the packet.
snmp-trap	Specifies the switch to send an <i>SNMP</i> trap when an event occurs.

Default

By default, DHCP snooping is disabled.

Usage Guidelines

Use this command to enable DHCP snooping on the switch.



Note

Snooping IP fragmented DHCP packets is not supported.

The violation action setting determines what action(s) the switch takes when a rouge DHCP server packet is seen on an untrusted port or the IP address of the originating server is not among those of

the configured trusted DHCP servers. The DHCP server packets are DHCP OFFER, ACK and NAK. The following list describes the violation actions:

- block-mac—The switch automatically generates an *ACL* to block the MAC address on that port. The switch does not blackhole that MAC address in the *FDB*. The switch can either temporarily or permanently block the MAC address.
- block-port—The switch blocks all incoming rogue DHCP packets on that port. The switch disables the port either temporarily or permanently to block the traffic on that port.
- none—The switch takes no action to drop the rogue DHCP packet or block the port, and so on. In this case, DHCP snooping continues to build and manage the DHCP bindings database and DHCP forwarding will continue in hardware as before.

Any violation that occurs causes the switch to generate an EMS log message. You can configure to suppress the log messages by configuring EMS log filters.

Displaying DHCP Snooping Information

To display the DHCP snooping configuration settings, use the following command:

```
show ip-security dhcp-snooping {vlan} vlan_name
```

To display the DHCP bindings database, use the following command:

```
show ip-security dhcp-snooping entries {vlan} vlan_name
```

To display any violations that occur, use the following command:

```
show ip-security dhcp-snooping violations {vlan} vlan_name
```

Example

The following example enables DHCP snooping on the switch and has the switch block DHCP packets from port 1:1:

```
enable ip-security dhcp-snooping vlan snoop ports 1:1 violation-action drop-packet block-port
```

History

This command was first available in ExtremeXOS 11.6.

Dynamic VLAN option added in ExtremeXOS 30.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable ip-security source-ip-lockdown ports

```
enable ip-security source-ip-lockdown ports [all | ports]
```

Description

Enables the source IP lockdown feature on one or more ports.

Syntax Description

all	Specifies all ports for which source IP lockdown should be enabled.
<i>ports</i>	Specifies one or more ports for which source IP lockdown should be enabled.

Default

By default, source IP lockdown is disabled on the switch.

Usage Guidelines



Note

Source-IP lockdown cannot be enabled on load sharing ports.

Source IP lockdown prevents IP address spoofing by automatically placing source IP address filters on specified ports. If configured, source IP lockdown allows only traffic from a valid *DHCP*-assigned address obtained by a DHCP snooping-enabled port or an authenticated static IP address to enter the network.

To configure source IP lockdown, you must enable DHCP snooping on the ports connected to the DHCP server and DHCP client before you enable source IP lockdown. You must enable source IP lockdown on the ports connected to the DHCP client, not on the ports connected to the DHCP server. The same DHCP bindings database created when you enable DHCP snooping is also used by the source IP lockdown feature to create ACLs that permit traffic from DHCP clients. All other traffic is dropped. In addition, the DHCP snooping violation action setting determines what action(s) the switch takes when a rogue DHCP server packet is seen on an untrusted port.

To enable DHCP snooping, use the following command:

```
enable ip-security dhcp-snooping {vlan} vlan_name ports [all | ports]  
violation-action [drop-packet {[block-mac | block-port] [duration  
duration_in_seconds | permanently] | none}}] {snmp-trap}
```

Displaying Source IP Lockdown Information

To display the source IP lockdown configuration on the switch, use the following command:

```
show ip-security source-ip-lockdown
```

Example

The following command enables source IP lockdown on ports 1:1 and 1:4:

```
enable ip-security source-ip-lockdown ports 1:1, 1:4
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable iqagent

```
enable iqagent
```

Description

Enables the ExtremeCloud™ IQ Agent.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

To disable the IQ Agent, use the command `disable iqagent`.

To view the state if the IQ Agent, use the command `show iqagent discovery` without the **discovery** option.

Example

The following example enables the IQ Agent:

```
# enables iqagent
```

History

This command was first available in ExtremeXOS 31.1.

Platform Availability

This command is available on the ExtremeSwitching 5520 platform.

enable irdp

```
enable irdp {vlan name}
```

Description

Enables the generation of *ICMP* router advertisement messages on one or all *VLANs*.

Syntax Description

<i>name</i>	Specifies a VLAN name.
-------------	------------------------

Default

Disabled.

Usage Guidelines

ICMP Router Discovery Protocol (IRDP) allows client machines to determine what default gateway address to use. The switch sends out IP packets at the specified intervals identifying itself as a default router. IRDP enabled client machines use this information to determine which gateway address to use for routing data packets to other networks.

If no optional argument is specified, all the IP interfaces are affected.

Example

The following example enables IRDP on VLAN "accounting":

```
enable irdp vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms that use the Edge, Advanced Edge, or Core license. For information on the licenses available for each platform, see the [Switch Engine 32.3 Feature License Requirements](#) document.[link-22.1"/>](#)

enable isis

```
enable isis {area area_name}
```

Description

This command enables the specified IS-IS router process on the current virtual router.

Syntax Description

<i>area_name</i>	Specifies the name of the IS-IS router process to be enabled.
------------------	---

Default

Disabled.

Usage Guidelines

If no area name is specified, all IS-IS router processes on the current virtual router are enabled. Once a router process is enabled, IS-IS PDUs are sent and processed provided that the following conditions are met:

- The router process has a system ID and area address configured.
- At least one associated *VLAN* interface has IPv4 or IPv6 forwarding enabled.

This command has no effect on router processes that are already enabled.

Example

The following command enables the IS-IS process named areax:

```
enable isis area areax
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Premier license as described in the *Switch Engine 32.3 Feature License Requirements* document.

enable isis area adjacency-check

```
enable isis area area_name adjacency-check {ipv4 | ipv6}
```

Description

This command enables the checking of the following TLVs when forming adjacencies: Protocols Supported and IP Interface Address.

Syntax Description

<i>area_name</i>	Specifies the name of the IS-IS router process that should perform the adjacency check.
ipv4	Specifies that the adjacency check is to be performed in IPv4 interfaces.
ipv6	Specifies that the adjacency check is to be performed in IPv6 interfaces.

Default

ipv4/ipv6: Enabled.

Usage Guidelines

When enabled for IPv4, IPv4 adjacencies may only be formed with neighbors whose connected interface supports IPv4 and is on the same subnet as the receiving interface. Similarly, when enabled for IPv6, IPv6 adjacencies may only be formed with neighbors whose connected interface supports IPv6 and is on the same link local subnet as the receiving interface. For each enabled protocol, if both criteria are not met, received Hello PDUs are discarded. By default, IPv4 routing is affected by this command. The optional `ipv6` keyword enables adjacency checking for IPv6 interfaces on the specified router process. It may be necessary to disable adjacency checking in multi-topology environments where a neighbor may only form an IPv4 or an IPv6 adjacency, but not both.

Example

The following command directs the IS-IS process named `areax` to perform adjacency checks on IPv6 interfaces:

```
enable isis area areax adjacency-check ipv6
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Premier license as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

enable isis area dynamic-hostname

`enable isis area area_name dynamic-hostname [area-name | snmp-name]`

Description

This command enables the dynamic hostname feature, which displays either the area name or the *SNMP* name instead of a IS-IS router system ID in select show commands.

Syntax Description

<i>area_name</i>	Specifies the name of the IS-IS process for which the dynamic-hostname feature is to be enabled.
area-name	Specifies that affected show commands display the area name instead of the IS-IS system ID.
snmp-name	Specifies that affected show commands display the SNMP name instead of the IS-IS system ID.

Default

Disabled.

Usage Guidelines

This command enables support for the dynamic hostname exchange feature defined by RFC 2763.

Example

The following command enables the display of IS-IS area names:

```
enable isis area areax dynamic-hostname area-name
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Premier license as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

enable isis area export

```
enable isis area area_name export {ipv4} route-type [policy | metric
mvalue {metric-type [internal | external]}} {level[1 | 2 | both-1-
and-2]}
```

Description

This command enables IPv4 route redistribution into IS-IS for direct, static, [BGP](#), [RIP](#), or [OSPF](#) routes.

Syntax Description

<i>area_name</i>	Specifies the IS-IS router process that receives the exported routes.
ipv4	Specifies that the redistributed routes are for use in IPv4 IS-IS routing.
<i>route-type</i>	Selects the type of route for export. The valid route types are: bgp, direct, e-bgp, i-bgp, ospf, ospf-extern1, ospf-extern2, ospf-inter, ospf-intra, rip, and static.
<i>policy</i>	Specifies a policy that controls how routes are redistributed into IS-IS.
<i>mvalue</i>	Specifies a metric to assign to the routes exported to IS-IS. The range is 0 to 4261412864.
metric-type [internal external]	Specifies a metric type, which is internal or external, to assign to the routes exported to IS-IS.
level [1 2 both-1-and-2]	Limits the use of redistributed routes to level 1, level 2, or both.

Default

All types are disabled.

Usage Guidelines

If wide metrics are enabled, redistributed routes are included in the Extended IP Reachability TLV in LSPs. If wide metrics are not enabled, redistributed routes are added to IP External Reachability TLV in LSPs. For policies, the nlri match attribute is supported, and the cost, cost-type internal, permit, and deny set attributes are supported.

Example

The following command exports RIP routes to IS-IS and assigns the internal metric type and metric value 5 to the redistributed routes:

```
enable isis area areax export rip metric 5 metric-type internal
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Premier license as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

enable isis area export ipv6

```
enable isis area area_name export ipv6 route-type [policy | metric mvalue] {level [1 | 2 | both-1-and-2] }
```

Description

This command enables IPv6 route redistribution into IS-IS for direct, static, [RIPng](#), or [OSPFv3](#) routes.

Syntax Description

<i>area_name</i>	Specifies the IS-IS router process that receives the exported routes.
<i>route-type</i>	Selects the type of route for export. The valid route types are: direct, ospfv3, ospfv3-extern1, ospfv3-extern2, ospfv3-inter, ospfv3-intra, ripng, bgp, and static.
<i>policy</i>	Specifies a policy that controls how routes are redistributed into IS-IS.
<i>mvalue</i>	Specifies a metric to assign to the routes exported to IS-IS. The range is 0 to 4261412864.
level [1 2 both-1-and-2]	Limits the use of redistributed routes to level 1, level 2, or both.

Default

All types are disabled.

Usage Guidelines

If a policy is specified, the policy is used to determine what specific routes are redistributed into IS-IS. Otherwise, the specified metric and type are assigned to the redistributed routes. Redistributed routes are added to the IPv6 External Reachability TLV in LSPs. For policies, the nlri match attribute is supported, and the cost, cost-type internal, permit, and deny set attributes are supported.

Example

The following command exports RIPng routes to IS-IS and assigns the internal metric type and metric value 5 to the redistributed routes:

```
enable isis area areax export ipv6 ripng metric 5 metric-type internal
```

History

This command was first available in ExtremeXOS 12.1.
Support for BGP was added in ExtremeXOS 12.6.0-BGP.

Platform Availability

This command is available on platforms with a Premier license as described in the Switch Engine 32.3 Feature License Requirements document.

enable isis area originate-default

```
enable isis area area_name originate-default {ipv4 | ipv6}
```

Description

This command causes the specified IS-IS router process to generate the default route in its LSPs.

Syntax Description

<i>area_name</i>	Specifies the name of the IS-IS router process that should generate the default route.
ipv4	Specifies that the router process should generate the default IPv4 route.
ipv6	Specifies that the router process should generate the default IPv6 route.

Default

IPv4: Disabled

IPv6: Disabled

Usage Guidelines

This applies to level 2 routing only. In contrast, level 1 routers compute the default route as the nearest attached L1/L2 router. When enabled, the router process generates an IPv4 default route unless the ipv6 option is specified. Only one level 2 router in the IS-IS domain should be configured to originate a default route. This command has no effect on router processes that are already enabled for default route origination or on router processes that are level 1-only.

Example

The following command directs the IS-IS process named areax to generate the default IPv4 route in it's LSPs:

```
enable isis area areax originate-default
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Premier license as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

enable isis area overload-bit

```
enable isis area area_name overload-bit {suppress [external | interlevel | all] }
```

Description

This command enables the overload-bit feature, which signals other routers that they are no longer permitted to use this router as a transit or forwarding node.

Syntax Description

<i>area_name</i>	Specifies the area name of the IS-IS process for which this feature is to be enabled.
suppress	Specifies that one or all types of reachability information is to be suppressed or excluded from LSPs.
external	When included with the suppress option, this specifies that external reachability information is to be excluded from LSPs.

interlevel	When included with the suppress option, this specifies that interlevel reachability information is to be excluded from LSPs.
all	When included with the suppress option, this specifies that external and interlevel reachability information is to be excluded from LSPs.

Default

Disabled.

Usage Guidelines

When the overload bit feature is enabled, the router process still receives and processes LSPs.

Example

The following command enables the overload bit feature for areax:

```
enable isis area areax overload-bit
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Premier license as described in the [**Switch Engine 32.3 Feature License Requirements**](#) document.

enable isis hello-padding

```
enable isis [vlan all | {vlan} vlan_name] hello-padding
```

Description

This command enables the padding of hello PDUs on one or all [VLANs](#).

Syntax Description

vlan all	Enables hello padding on all IS-IS VLANs.
<i>vlan_name</i>	Specifies a single VLAN on which hello padding is enabled.

Default

Enabled.

Usage Guidelines

When hello padding is enabled, IS-IS pads hello packets to the interface MTU. This is used among neighbors to verify that adjacencies have the same MTU configured on either end. The disadvantage of hello padding is the price of bandwidth consumed by larger packets.

Example

The following command enables hello padding on the SJvlan VLAN:

```
enable isis SJvlan hello-padding
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Premier license as described in the [**Switch Engine 32.3 Feature License Requirements**](#) document.

enable isis restart-helper

```
enable isis restart-helper
```

Description

This command enables the IS-IS router to act as a restart helper according to draft-ietf-isis-restart-02—Restart signaling for IS-IS.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

None.

Example

The following command enables the IS-IS restart helper:

```
enable isis restart-helper
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Premier license as described in the [***Switch Engine 32.3 Feature License Requirements***](#) document.

enable jumbo-frame ports

```
enable jumbo-frame ports [all | port_list]
```

Description

Enables support on the physical ports that will carry jumbo frames.

Syntax Description

all	Specifies ports.
<i>port_list</i>	Specifies one or more slots and ports.

Default

Disabled.

Usage Guidelines

Increases performance to back-end servers or allows for VMAN 802.1Q encapsulations.

You can configure the maximum size of a jumbo frame if you want to use a different size than the default value of 9216. Use the [`configure jumbo-frame-size`](#) command to configure the size.

This setting is preserved across reboots.

Example

The following command enables jumbo frame support on a switch:

```
enable jumbo-frame ports all
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable l2vpn

```
enable l2vpn [vppls [vppls_name | all] | vpws [vpws_name | all]]
```

Description

Enables the specified VPLS or VPWS.

Syntax Description

<i>vppls_name</i>	Identifies the VPLS within the switch (character string).
<i>vpws_name</i>	Identifies the VPWS within the switch (character string).
all	Specifies all VPLS or VPWS instances.

Default

All newly created VPLS or VPWS instances are enabled.

Usage Guidelines

When enabled, VPLS or VPWS attempts to establish sessions between all configured peers. Services must be configured and enabled for sessions to be established successfully.

The **l2vpn** keyword is introduced in ExtremeXOS Release 12.4 and is required when enabling a VPWS. For backward compatibility, the **l2vpn** keyword is optional when enabling a VPLS. However, this keyword will be required in a future release, so we recommend that you use this keyword for new configurations and scripts.

Example

The following command enables the VPLS instance myvppls:

```
enable vppls myvppls
```

The following command enables the VPWS instance myvpws:

```
enable l2vpn vpws myvpws
```

History

This command was first available in ExtremeXOS 11.6.

The **l2vpn** and **vpws** keywords were first available in ExtremeXOS 12.4.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

enable l2vpn health-check vccv

```
enable l2vpn [vpls vpls_name | vpws vpws_name] health-check vccv
```

Description

Enables the VCCV health check feature on the specified VPLS or VPWS.

Syntax Description

vpls_name	Identifies the VPLS for which health check is to be enabled.
vpws_name	Identifies the VPWS for which health check is to be enabled.

Default

Health check is disabled.

Usage Guidelines

Health check must be enabled on both ends of a PW to verify connectivity between two VPLS or VPWS peers. Both VCCV peers negotiate capabilities at PW setup. A single VCCV session monitors a single PW. Therefore, a VPLS with multiple PWs will have multiple VCCV sessions to multiple peers.

VCCV in ExtremeXOS uses LSP ping to verify connectivity.

The **l2vpn** keyword was introduced in ExtremeXOS Release 12.4 and is required when enabling health check for a VPWS instance. For backward compatibility, the **l2vpn** keyword is optional when enabling health check for a VPLS instance. However, this keyword will be required in a future release, so we recommend that you use this keyword for new configurations and scripts.

Example

The following command enables the health check feature on the VPLS instance myvpls:

```
enable l2vpn vpls myvpls health-check vccv
```

History

This command was first available in ExtremeXOS 12.1.

The **l2vpn** and **vpws** keywords were first available in ExtremeXOS 12.4.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

enable l2vpn service

```
enable l2vpn [vppls [vppls_name | all] | vpws [vpws_name | all]] service
```

Description

Enables the configured services for the specified VPLS or VPWS.

Syntax Description

<i>vppls_name</i>	Identifies the VPLS within the switch (character string).
<i>vpws_name</i>	Identifies the VPWS within the switch (character string).
all	Specifies all VPLS or VPWS instances.

Default

Enabled.

Usage Guidelines

When services are disabled, the VPLS or VPWS is withdrawn from all peer sessions. The keyword **all** enables services for all VPLS or VPWS instances.

The **l2vpn** keyword was introduced in ExtremeXOS Release 12.4 and is required when enabling services for a VPWS instance. For backward compatibility, the **l2vpn** keyword is optional when enabling services for a VPLS instance. However, this keyword will be required in a future release, so we recommend that you use this keyword for new configurations and scripts.

Example

The following command enables the configured VPLS services for the specified VPLS instance:

```
enable l2vpn vppls myvppls service
```

History

This command was first available in ExtremeXOS 11.6.

The **l2vpn** and **vpws** keywords were first available in ExtremeXOS 12.4.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

enable l2vpn sharing

```
enable l2vpn sharing
```

Description

Enables LSP sharing for Layer 2 VPN pseudowires .

Syntax Description

This command has no keywords or arguments.

Default

Disabled.

Usage Guidelines

This command enables LSP sharing for L2VPN PWs. When LSP sharing is enabled, up to 16 named LSPs are used for a PW. When LSP sharing is disabled, only 1 named LSP is used for a PW.

If LSP Sharing is disabled, and more than 1 Transport LSP is programmed into HW, all but 1 Transport LSP is removed from HW, and the configuration is preserved. If LSP Sharing is enabled, and more than 1 Transport LSP was previously configured, the remaining LSPs is programmed into HW as they become available for use.

Example

The following command enables LSP sharing for L2VPN PWs:

```
enable l2vpn sharing
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

enable l2vpn vpls peer fdb send-mac-withdrawal

```
enable l2vpn vpls peer [ipaddress | all] fdb send-mac-withdrawal
```

Description

Enables the Layer 2 VPN MAC address withdrawal capability.

Syntax Description

l2vpn	Designates L2 VPN configuration.
vpls	Designates VPLS of MPLS configuration.
peer	Designates VPLS peer.
<i>ipaddress</i>	Selects the VPLS peer of the provided IP address.
all	Selects all VPLS peers.
fdb	Designates FDB.
send-mac-withdrawal	Enables sending the MAC address withdrawal message.

Default

Enabled.

Usage Guidelines

Use this command to enable *FDB* MAC withdrawal after it has been disabled.

Example

The following command enables MAC address withdrawal message for all VPLS peers:

```
# enable l2vpn vpls peer all fdb send-mac-withdrawal
```

History

This command was first available in ExtremeXOS 12.1.

The **l2vpn** keyword was added in ExtremeXOS 12.4.

Platform Availability

This command is available only on the platforms that support *MPLS* as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

enable learning iparp sender-mac

```
enable learning iparp {request | reply | both-request-and-reply} {vr  
vr_name} sender-mac
```

Description

Enables MAC address learning from the payload of IP ARP packets.

Syntax Description

request	Enables learning only for IP ARP request packets.
reply	Enables learning only for IP ARP reply packets.
both-request-and-reply	Enables learning for both request and reply packets.
<i>vr_name</i>	Specifies a virtual router.

Default

Disabled.

Usage Guidelines

To view the configuration for this feature, use the following command: `show iparp`

Example

The following command enables MAC address learning from the payload of reply IP ARP packets:

```
enable learning iparp reply sender-mac
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable learning port

```
enable learning {drop-packets} ports [all | port_list]
```

Description

Enables MAC address learning on one or more ports.

Syntax Description

drop-packets	Forwards <i>EDP</i> packets, and drops all unicast, multicast, and broadcast packets from a source address not in the <i>FDB</i> . No further processing occurs for dropped packets.
all	Specifies all ports.
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

Enabled.

Usage Guidelines

Use this command to enable MAC address learning on one or more ports.

Example

The following example enables MAC address learning on ports 7 and 8 on a switch:

```
enable learning ports 7-8
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable led locator

```
enable led locator {timeout [seconds | none]} {pattern [alternating |  
flash-all | high-to-low | scanner]} {slot [ slot | all ]}
```

Description

Configures the front panel LEDs to flash so a switch can be easily located in a crowded lab/data center.

Syntax Description

timeout	Limit the LED display time to <i>seconds</i> before returning to normal operation.
<i>seconds</i>	The length of time to display the flashing LEDs. The default is 300 seconds. The maximum value is 1 week (604800 seconds).
none	Display LED pattern until disabled.
pattern	Configures the LED display pattern.
alternating	Groups of LEDs are lit in alternating patterns (Default).
flash-all	All LEDs flash on and off.
high-to-low	LED's are lit in descending port order.
scanner	A group of 4 LED's is lit back and forth.
slot <i>slot</i>	Slot number.
all	All slots.

Default

The default **timeout** length is 300 seconds.

The default pattern is alternating.

Usage Guidelines

Use this command to enable the front panel LEDs to flash so that a switch can be easily located in a crowded lab, or data center.

Bridge Port Extenders (BPEs)

The LED locator service works for both directly attached and cascaded bridge port extenders (BPEs). Only the default LED pattern of alternating is supported by the BPEs. This causes alternating flashing of groups of 8 port LEDs.

Example

The following example enables the front panel LEDs to flash in an alternating pattern for one hour on all slots:

```
enable led locator timeout 3600 pattern alternating all
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable license

```
enable license {software} [key ]
```

Description

Enables software license or feature pack that allows you to use advanced features.

Syntax Description

software	Applies base license.
<i>key</i>	Specifies your hexadecimal license key in format xxxx-xxxx-xxxx-xxxx-xxxx (10 hex digits) or xxxx-xxxx-xxxx-xxxx-xxxx-xxxx-xxxx (14 hex digits).

Default

N/A.

Usage Guidelines

The software license levels that apply to ExtremeXOS software are described in the [Switch Engine 32.3 Feature License Requirements](#) document.

To obtain a software license, specify the key in the format xxxx-xxxx-xxxx-xxxx-xxxx.

You obtain the software license key (or feature pack key) either by ordering it from the factory or by obtaining a license voucher from your Extreme Networks supplier. You can obtain a regular software license or a trial software license, which allows you use of the license for either 30, 60 or 90 days; you cannot downgrade software licenses.

The voucher contains all the necessary information on the software license, whether regular or trial, and number of days for trial software license.

After you enable the software license or feature pack by entering the software key, the system returns a message that you either successfully or unsuccessfully set the license.

Once you enable the software license (or if you do not use the correct key, attempt to downgrade the license, or already installed the software license) you see one of the following messages:

```
Enabled license successfully. Error: Unable to set license using supplied key. Error:
Unable to set license - downgrade of licenses is not supported. Error: Unable to set
license - license is already enabled. Error: Unable to set license - trial license already
enabled.
```

If you enable a trial license, the system generates a daily message showing the number of days until expiry.

If you attempt to execute a command and you do not either have the required software license or have reached the limits defined by the current software license level, the system returns one of the following messages:

```
Error: This command cannot be executed at the current license level. Error: You have
reached the maximum limit for this feature at this license level.
```

If you attempt to execute a command and you do not have the required feature pack, the system also returns a message.

To protect against attacks to install maliciously created license keys, the system has an exponential delay of each failed attempt to install a license.

To view the type of software license you are currently running on the switch, use the [show licenses](#) command. The license key number is not displayed, but the type of software license is displayed in the [show licenses](#) output. This command can be run on any node in a SummitStack, regardless of its node role (master, standby, or backup).

Messages for different scenarios:

- Key format is incorrect:
"Error: Incorrect key format."
- Attempted to apply Switch Port Speed License to another switch model:

```
Error: Unable to set license - platform incompatible with license.
```

- Attempted to apply Switch Port Speed License beyond number of ports available on switch:

```
Error: Unable to set license - platform only supports ports in range 1 to <max for platform>.
```
- EEPROM Read/Write failure for Switch Port Speed License:

```
Error: Unable to set license. Read from EEPROM failed.
```
- Attempted to apply Switch Port Speed License when it is already applied:

```
Error: Unable to set license - license is already enabled.
```
- Attempted to apply Switch Port Speed License license for fewer ports groups than is currently enabled. You cannot downgrade the license this way. However, you can remove the license using the **clear license** command, and then apply a license enabling fewer port groups:

```
Error: Unable to set license - downgrade of port speed license not supported.  
<num_ports> ports already licensed. Current license can be cleared via 'clear license-  
info port-speed'.
```
- EEPROM Read/Write failure for Switch Port Speed License:

```
Error: Unable to set license - write to EEPROM failed.
```

Example

The following command enables a software license on the switch:

```
enable license 2d5e-0e84-e87d-c3fe-bfff
```

```
Warning: A reboot switch or disable and enable slot 3 is required before the new license  
takes effect.
```

History

This command was first available in ExtremeXOS 11.1.

The software parameter was added in ExtremeXOS 11.6.

The *capacity-key* variable was added in ExtremeXOS 15.4.

The *capacity-key* variable was removed in ExtremeXOS 22.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable license file

```
enable license file filename
```

Description

Enables the text file that applies software licenses and feature packs licenses to more than one switch at a time.

Syntax Description

<i>filename</i>	Specifies the file name that you download onto the switch using TFTP; the file extension is .xlic.
-----------------	--

Default

N/A.

Usage Guidelines

You download the license file to the switch using TFTP or SCP. The file name extension for this file is *xlic*; for example, you may see a file named *systemlic.xlic*.

Using this file, you enable the software and feature pack licenses for more than one switch simultaneously. The file can contain licenses for some or all of the Extreme Networks switches that the customer owns. During upload, only those license keys destined for the specific switch are used to attempt enabling the licenses. The license file is a text file that has the switch serial number, software license type, and license key; it is removed from the switch after the licenses are enabled.

After you enable the license file, the system returns one or more of the following messages:

```
Enabled license successfully. Error: Unable to set license
<license_name> using supplied key. Error: Unable to set license
<license_name> - downgrade of licenses is not supported. Error: Unable
to set license <license_name> - license is already enabled. Error:
Unable to set license <license_name> - trial license already enabled.
```

To protect against attacks to install maliciously created license keys, the system has an exponential delay of each failed attempt to install a license.

Example

The following command enables a license file on the specified Extreme Networks switches:

```
enable license file santaclara.xlic
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable lldp ports

```
enable lldp ports [all | port_list] {receive-only | transmit-only}
```


Description

Enables *LLDP* transmit mode, receive mode, or transmit and receive mode. If the transmit-only or receive-only option is not specified, both transmit and receive modes are enabled.

Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.
receive-only	Specifies that the port only receives LLDP messages.
transmit-only	Specifies that the port only transmits LLDP messages.

Default

Enabled.

Usage Guidelines

If you do not specify an option, the port is enabled to both transmit and receive LLDP messages.

Once the port is enabled for LLDP in one mode and you issue another `enable lldp ports` command for another mode, that second mode replaces the original mode. For example, you might originally enable several ports to only receive LLDP messages and then want those ports to both receive and transmit LLDP messages. In that case, you issue the `enable lldp ports` command with no variables (and the receive-and-transmit mode replaces the receive-only mode).

To verify the port setting for LLDP, use the `show lldp {port [all |port_list]} {detailed}` command.

Example

The following example enables LLDP transmit and receive mode on port 1:4.

```
enable lldp port 1:4
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable log debug-mode

```
enable log debug-mode
```

Description

Enables debug mode. The switch generates debug events.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command enables debug mode. Debug mode must be enabled prior to configuring advanced debugging capabilities. These include allowing debug messages, which can severely degrade performance. For typical network device monitoring, debug mode should remain disabled, the default setting. Debug mode should only be enabled when advised by technical support, or when advanced diagnosis is required. The debug mode setting is saved to FLASH.

The following configuration options require that debug mode be enabled:

- Including a severity of debug-summary, debug-verbose, or debug-data when configuring filters.
- Target format options process-name, process-id, source-function, and source-line.

Example

The following command enables debug mode:

```
enable log debug-mode
```

When you enable debug mode, the following message appears:

```
WARNING: Debug mode should only be enabled when advised by technical
support, or when advanced diagnosis is required. Performance degradation
is possible. Debug mode now enabled.
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable log display

```
enable log display
```

Description

Enables a running real-time display of log messages on the console display. In a stack, this command is applicable only to Master and Backup nodes. You cannot run this command on standby nodes.

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).

You configure the messages displayed in the log using the `configure log display`, or configure log target console-display commands.

You can also use this command to control logging to different targets. This command is equivalent to `enable log target console-display` command.

To change the log filter association, severity threshold, or match expression for messages sent to the console display, use the configure log target console-display command

Example

The following command enables a real-time display of log messages:

```
enable log display
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable log target

```
enable log target [console | memory-buffer | nvram | primary-
node|backup-node| session | syslog [all | ipaddress udp-port
{udp_port} | ipPort | ipaddress tls_port {tls_port}] {vr vr_name}
{local0...local17}]]
```

Description

Starts sending log messages to the specified target.

Syntax Description

console	Specifies the console display.
memory-buffer	Specifies the switch memory buffer.
nvr am	Specifies the switch NVRAM.
primary-node	Specifies the primary node of a stack.
backup-node	Specifies the backup node of a stack.
session	Specifies the current session (including console display).
syslog	Specifies a syslog target.
all	Specifies all of the remote syslog servers.
<i>ipaddress</i>	Specifies the syslog IP address.
<i>ipPort</i>	Specifies the UDP port number for the syslog target.
tls_port	Specifies remote Syslog server Transport Layer Security (TLS) for connection type.
<i>tls_port</i>	TLS port number.
<i>vr_name</i>	Specifies the virtual router that can reach the server IP address. Note: User-created VRs are supported only on the platforms listed for this feature in the Switch Engine 32.3 Feature License Requirements document.
local0 ... local7	Specifies the local syslog facility.

Default

Enabled, for memory buffer and NVRAM; all other targets are disabled by default.

Usage Guidelines

This command starts sending messages to the specified target. By default, the memory-buffer, NVRAM, primary node, and backup node targets are enabled. Other targets must be enabled before messages are sent to those targets.

Configuration changes to the session target are in effect only for the duration of the console display or Telnet session, and are not saved in FLASH. Others are saved in FLASH.

You can also use the following command to enable displaying the log on the console: `enable log display`

This command is equivalent to the `enable log target console-display` command.

Example

The following example enables log messages on the current session:

```
enable log target session
```

History

This command was first available in ExtremeXOS 10.1.

The **ipPort** parameter was first available in ExtremeXOS 11.0.

The **udp-port** parameter was added in ExtremeXOS 21.1.

Transport Layer Security (TLS) option added in ExtremeXOS 22.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable log target upm

```
enable log target upm {upm_profile_name}
```

Description

Enables the specified UPM log target.

Syntax Description

<i>upm_profile_name</i>	Specifies the name of the UPM log target to be enabled.
-------------------------	---

Default

N/A.

Usage Guidelines

UPM log targets are disabled when they are created.

Example

The following command enables the UPM log target testprofile1:

```
enable log target upm testprofile1
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the Universal Port feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

enable log target xml-notification

```
enable log target xml-notification xml_target_name
```

Description

Enables a Web server target.

Syntax Description

<i>xml_target_name</i>	Specifies the name of the xml-notification target.
------------------------	--

Default

N/A.

Usage Guidelines

Use this command to enable a web server target for EMS.

Example

The following command enables the web server target target2:

```
enable log target xml-notification target2
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable loopback-mode vlan

```
enable loopback-mode vlan [vlan_name | vlan_list]
```

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

<i>vlan_name</i>	Specifies a VLAN name.
<i>vlan_list</i>	Specifies a VLAN list of IDs.

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 example 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 ExtremeXOS 10.1.

The *vlan_list* option was added in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable mac-lockdown-timeout ports

```
enable mac-lockdown-timeout ports [all | port_list]
```

Description

Enables the MAC address lock down timeout feature for the specified port or group of ports or for all ports on the switch.

Syntax Description

all	Specifies all ports.
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

By default, the MAC address lock down timeout feature is disabled.

Usage Guidelines

You cannot enable the MAC lock down timer on a port that also has the lock learning feature enabled.

Example

The following command enables the MAC address lock down timeout feature for ports 2:3, 2:4, and 2:6:

```
enable mac-lockdown-timeout ports 2:3, 2:4, 2:6
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable mac-locking ports

```
enable mac-locking ports [port_list | all]
```

Description

Enables MAC locking on the specified port.

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
all	Specifies all ports.

Default

MAC locking is disabled by default.

Usage Guidelines

To enable MAC locking on a specific port, you must enable MAC locking on the switch and on the port. Use the `enable mac-locking` command to enable MAC locking on the switch.

You cannot enable MAC locking on a port if limit-learning or lock-learning is configured on the port for any *VLAN*.

Example

The following example enables MAC locking on port 14.

```
enable mac-locking ports 14
```


History

This command was first available in ExtremeXOS 15.7.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable mac-locking

```
enable mac-locking
```

Description

Enables MAC locking globally on the switch.

Syntax Description

This command has no arguments or variables.

Default

MAC locking is disabled by default.

Usage Guidelines

To enable MAC locking on a specific port, you must enable MAC locking on the switch and on the port. Use the [enable mac-locking ports](#) command to enable MAC locking on a port.

Example

The following example enables MAC locking on the switch.

```
enable mac-locking
```

History

This command was first available in ExtremeXOS 15.7.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable mirror

```
enable mirror mirror_name
```

Description

Enables a mirror instance.

Syntax Description

<i>mirror_name</i>	Specifies the mirror name.
--------------------	----------------------------

Default

Disabled.

Usage Guidelines

Use this command to enable a mirror instance. An instance may be enabled without source filters defined (per current function), but no traffic will be mirrored until source filters are added.

Example

The following example enables a mirror instance named "mirror1" :

```
enable mirror mirror1
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable mirror control_index

```
enable mirror control_index {mirror mirror_name}
```

Description

Enables a Mirror MIB instance or the assigned instance to an existing mirror.

Syntax Description

<i>control_index</i>	Selects the Mirror MIB instance to enable. Range is 1 through 4.
mirror	Designates specifying a mirror name associated within the specified control index.
<i>mirror_name</i>	Specifies the mirror name associated within the specified control index.

Default

Disabled.

Usage Guidelines

Specifying a mirror name only enables that mirror within the Mirror MIB group (control index).

Example

The following example enables Mirror MIB specified by control index "1":

```
# enable mirror 1
```

The following example enables the mirror named "m1" within the Mirror MIB specified by control index "1":

```
# enable 1 mirror m1
```

History

This command was first available in ExtremeXOS 30.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable mirror to port

```
enable mirror to [port port | port-list port_list loopback-port port]  
                {remote-tag tag}
```

Description

Dedicates a port on the switch to be the mirror output port, or the monitor port.

Syntax Description

<i>port</i>	Specifies the mirror output port.
<i>port_list</i>	Specifies the list of ports where traffic is to be mirrored.
loopback-port	Specifies an otherwise unused port required when mirroring to a <i>port_list</i> . The loopback-port is not available for switching user data traffic.
<i>port</i>	Specifies a single loopback port that is used internally to provide this feature.
remote-tag	Specifies the value of the <u>VLAN</u> ID used by the mirrored packets when egressing the monitor port.

Default

Disabled.

Usage Guidelines

Port mirroring configures the switch to copy all traffic associated with one or more ports, VLANs or virtual ports. A virtual port is a combination of a VLAN and a port. The monitor port(s) 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.

Up to 16 mirroring filters and up to four monitor ports 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.

You cannot run ELSM and mirroring on the same port. If you attempt to enable mirroring on a port that is already enabled for ELSM, the switch returns a message similar to the following:

Error: Port mirroring cannot be enabled on an ELSM enabled port.

Standalone Switches and SummitStacks

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(s). You can specify which traffic the port mirrors:
 - Ingress—Mirrors traffic received at the port.
 - Egress—Mirrors traffic sent from the port.
 - Ingress and egress—Mirrors traffic either received at the port or sent from the port.

(If you omit the optional parameters, all traffic is forwarded; the default for port-based mirroring is ingress and egress).

- **VLAN**—All data to 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.
- ExtremeSwitching series switches support a maximum of 128 mirroring filters with the restriction that a maximum of 16 VLAN and/or virtual port (port + VLAN) filters may be configured.
- ExtremeXOS supports up to 16 monitor ports for one-to-many mirroring.
- Only traffic ingressing a VLAN can be monitored; you cannot specify ingressing or egressing traffic when mirroring VLAN traffic.
- Ingress traffic is mirrored as it is received (on the wire).
- Packets which match both an ingress filter and an egress filter will result in two packets egressing the monitor port or ports.
- In normal mirroring, a monitor port cannot be added to a load share group. In one-to-many mirroring, a monitor port list can be added to a load share group, but a loopback port cannot be used in a load share group.
- You can run mirroring and sFlow on the same device when you are running ExtremeSwitching series switches.

- With a monitor port or ports on ExtremeSwitching series switches, all traffic ingressing the monitor port or ports is tagged only if the ingress packet is tagged. If the packet arrived at the ingress port as untagged, the packet egress the monitor port or ports as untagged.
- Two packets are mirrored when a packet encounters both an ingress and egress mirroring filter.
- The configuration of remote-tag does not require the creation of a VLAN with the same tag; on these platforms the existence of a VLAN with the same tag as a configured remote-tag is prevented. This combination is allowed so that an intermediate remote mirroring switch can configure remote mirroring using the same remote mirroring tag as other source switches in the network. Make sure that VLANs meant to carry normal user traffic are not configured with a tag used for remote mirroring.

When a VLAN is created with remote-tag, that tag is locked and a normal VLAN cannot have that tag. The tag is unique across the switch. Similarly if you try to create a remote-tag VLAN where remote-tag already exists in a normal VLAN as a VLAN tag, you cannot use that tag and the VLAN creation fails.

SummitStack Only

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(s). You can specify which traffic the port mirrors:
 - Ingress—Mirrors traffic received at the port.
 - Egress—Mirrors traffic sent from the port.
 - Ingress and egress—Mirrors traffic either received at the port or sent from the port.

(If you omit the optional parameters, all traffic is forwarded; the default for port-based mirroring is ingress and egress).

- **VLAN**—All data to 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.
- SummitStack supports a maximum of 128 mirroring filters with the restriction that a maximum of 16 VLAN and/or virtual port (port + VLAN) filters may be configured.
- ExtremeXOS supports up to 16 monitor ports for one-to-many mirroring.
- Only traffic ingressing a VLAN can be monitored; you cannot specify ingressing or egressing traffic when mirroring VLAN traffic.
- Ingress traffic is mirrored as it is received (on the wire).
- Two packets are mirrored when a packet encounters both an ingress and egress mirroring filter.
- When traffic is modified by hardware on egress, egress mirrored packets may not be transmitted out of the monitor port as they egressed the port containing the egress mirroring filter. For example, an egress mirrored packet that undergoes VLAN translation is mirrored with the untranslated VLAN ID. In addition, IP multicast packets which are egress mirrored contain the source MAC address and VLAN ID of the unmodified packet.
- You cannot include the monitor port for a SummitStack in a load-sharing group.
- You can run mirroring and sFlow on the same device when you are running a SummitStack.

- With a monitor port or ports, the mirrored packet is tagged only if the ingress packet is tagged (regardless of what module the ingressing port is on). If the packet arrived at the ingress port as untagged, the packet egress the monitor port(s) as untagged.
- You may see a packet mirrored twice. This occurs only if both the ingress mirrored port and the monitor port or ports are on the same one-half of the module and the egress mirrored port is either on the other one-half of that module or on another module.
- When traffic is modified by hardware on egress, egress mirrored packets may not be transmitted out of the monitor port as they egressed the port containing the egress mirroring filter. For example, an egress mirrored packet that undergoes VLAN translation is mirrored with the untranslated VLAN ID. In addition, IP multicast packets which are egress mirrored contain the source MAC address and VLAN ID of the unmodified packet.
- The configuration of remote-tag does not require the creation of a VLAN with the same tag; on these platforms the existence of a VLAN with the same tag as a configured remote-tag is prevented. This combination is allowed so that an intermediate remote mirroring switch can configure remote mirroring using the same remote mirroring tag as other source switches in the network. Make sure that VLANs meant to carry normal user traffic are not configured with a tag used for remote mirroring.
- When a VLAN is created with remote-tag, that tag is locked and a normal VLAN cannot have that tag. The tag is unique across the switch. Similarly if you try to create a remote-tag VLAN where remote-tag already exists in a normal VLAN as a VLAN tag, you cannot use that tag and the VLAN creation fails.

Example

The following example selects port 4 as the mirror, or monitor, port:

```
# enable mirror to port 4
```

History

This command was added in ExtremeXOS 15.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable mirror to remote-ip

```
enable mirror {mirror_name} to remote-ip remote_ip_address {{vr}
vr_name} {priority priority_value} {from [source_ip_address | auto-
source-ip]} {ping-check [on | off]}
```

Description

Enables traffic to be mirrored to the specified remote IPv4 destination address encapsulated in a GRE tunneled packet.

Syntax Description

<i>mirror_name</i>	Specifies the mirror instance name.
remote-ip	Specifies to send mirrored packets to specified destination remote IP address.
<i>remote_ip_address</i>	Specifies the remote destination IP address for mirrored packets.
vr	Specifies a virtual router of the remote IP address.
<i>vr_name</i>	Specifies the virtual router name. If not specified, VR of current command context is used.
from	Configures source IP address of encapsulated mirrored packets.
<i>source_ip_address</i>	Specifies the local source IPv4 address for encapsulated mirrored packets.
auto-source-ip	Automatically use source IP address of egress VLAN to be used to reach remote IP address.
ping-check	Configure ping health check for remote IP address.
on	Only send mirrored packets to remote IP address if periodic pings to remote IP address are successful (default).
off	Send mirrored packets to remote IP address without any ping health check, assuming MAC address and port of next hop IP address are static or learned.
priority	Configures a unique priority value for each redundant remote IP address of a mirror instance.
<i>priority_value</i>	Sets the unique priority value for the remote IP address. The priority value must be unique for each remote IP address in the mirror instance. The range is from 1 (least preferred) to 100 (most preferred). The default is 50.

Default

If a VR is not specified, the VR of the current command context is used.

Ping health check of the remote IP address is enabled unless otherwise specified.

The default priority value is 50.

Usage Guidelines

This command enables hardware mirroring of Ethernet frames to a specified remote IPv4 address, which can reside zero or more router hops away. This is useful for ExtremeAnalytics Application Telemetry or other forms of remote network analysis or monitoring.

Port mirroring configures the switch to copy all traffic associated with one or more ports, VLANs or virtual ports. A virtual port is a combination of a VLAN and a port. The monitor port(s) 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.

Up to 16 mirroring filters and up to four monitor ports 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.

You cannot run ELSM and mirroring on the same port. If you attempt to enable mirroring on a port that is already enabled for ELSM, the switch returns a message similar to the following:

```
Error: Port mirroring cannot be enabled on an ELSM enabled port.
```

For high availability, you can add up to four redundant remote IP addresses. When creating a mirror with this command, you can add one IP address. To add additional remote IP addresses, use the `configure mirror` *mirror_name* {**to** [**port** *port* | **port-list** *port_list* | **loopback** *port* *port*] | **remote-ip** {**add**} *remote_ip_address* [{**vr**} *vr_name*] {**from** [*source_ip_address* | **auto-source-ip**] } {**ping-check** [**on** | **off**] } } {**remote-tag** *rtag* | **port** **none**} {**priority** *priority_value*}**command**.

Standalone Switches and SummitStacks

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(s). You can specify which traffic the port mirrors:
 - Ingress—Mirrors traffic received at the port.
 - Egress—Mirrors traffic sent from the port.
 - Ingress and egress—Mirrors traffic either received at the port or sent from the port.

(If you omit the optional parameters, all traffic is forwarded; the default for port-based mirroring is ingress and egress).
- **VLAN**—All data to 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.
- ExtremeSwitching series switches support a maximum of 128 mirroring filters with the restriction that a maximum of 16 VLAN and/or virtual port (port + VLAN) filters may be configured.
- ExtremeXOS supports up to 16 monitor ports for one-to-many mirroring.
- Only traffic ingressing a VLAN can be monitored; you cannot specify ingressing or egressing traffic when mirroring VLAN traffic.
- Ingress traffic is mirrored as it is received (on the wire).
- Packets which match both an ingress filter and an egress filter will result in two packets egressing the monitor port or ports.
- In normal mirroring, a monitor port cannot be added to a load share group. In one-to-many mirroring, a monitor port list can be added to a load share group, but a loopback port cannot be used in a load share group.
- You can run mirroring and sFlow on the same device when you are running ExtremeSwitching series switches.
- With a monitor port or ports on ExtremeSwitching series switches, all traffic ingressing the monitor port or ports is tagged only if the ingress packet is tagged. If the packet arrived at the ingress port as untagged, the packet egress the monitor port or ports as untagged.
- Two packets are mirrored when a packet encounters both an ingress and egress mirroring filter.

SummitStack Only

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(s). You can specify which traffic the port mirrors:
 - Ingress—Mirrors traffic received at the port.
 - Egress—Mirrors traffic sent from the port.
 - Ingress and egress—Mirrors traffic either received at the port or sent from the port.

(If you omit the optional parameters, all traffic is forwarded; the default for port-based mirroring is ingress and egress).

- **VLAN**—All data to 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.
- SummitStack supports a maximum of 128 mirroring filters with the restriction that a maximum of 16 VLAN and/or virtual port (port + VLAN) filters may be configured.
- ExtremeXOS supports up to 16 monitor ports for one-to-many mirroring.
- Only traffic ingressing a VLAN can be monitored; you cannot specify ingressing or egressing traffic when mirroring VLAN traffic.
- Ingress traffic is mirrored as it is received (on the wire).
- Two packets are mirrored when a packet encounters both an ingress and egress mirroring filter.
- When traffic is modified by hardware on egress, egress mirrored packets may not be transmitted out of the monitor port as they egressed the port containing the egress mirroring filter. For example, an egress mirrored packet that undergoes VLAN translation is mirrored with the untranslated VLAN ID. In addition, IP multicast packets which are egress mirrored contain the source MAC address and VLAN ID of the unmodified packet.
- You cannot include the monitor port for a SummitStack in a load-sharing group.
- You can run mirroring and sFlow on the same device when you are running a SummitStack.
- With a monitor port or ports, the mirrored packet is tagged only if the ingress packet is tagged (regardless of what module the ingressing port is on). If the packet arrived at the ingress port as untagged, the packet egress the monitor port(s) as untagged.
- You may see a packet mirrored twice. This occurs only if both the ingress mirrored port and the monitor port or ports are on the same one-half of the module and the egress mirrored port is either on the other one-half of that module or on another module.
- When traffic is modified by hardware on egress, egress mirrored packets may not be transmitted out of the monitor port as they egressed the port containing the egress mirroring filter. For example, an egress mirrored packet that undergoes VLAN translation is mirrored with the untranslated VLAN ID. In addition, IP multicast packets which are egress mirrored contain the source MAC address and VLAN ID of the unmodified packet.
- The configuration of remote-tag does not require the creation of a VLAN with the same tag; on these platforms the existence of a VLAN with the same tag as a configured remote-tag is prevented. This combination is allowed so that an intermediate remote mirroring switch can configure remote mirroring using the same remote mirroring tag as other source switches in the network. Make sure

that VLANs meant to carry normal user traffic are not configured with a tag used for remote mirroring.

- When a VLAN is created with remote-tag, that tag is locked and a normal VLAN cannot have that tag. The tag is unique across the switch. Similarly if you try to create a remote-tag VLAN where remote-tag already exists in a normal VLAN as a VLAN tag, you cannot use that tag and the VLAN creation fails.

Example

The following example enables a mirroring instance named "analytics_chicago_1" to mirror packets to the remote IP address 1.2.3.4 with ping health check (default behavior) being performed on the remote IP address:

```
enable mirror analytics_chicago_1 to remote-ip 1.2.3.4
```

The following example enables a mirroring instance named "analytics_seattle_2" to mirror packets to the remote IP address 5.6.7.8 from the source IP address 10.1.1.1 without ping health check being performed on the remote IP address:

```
enable mirror analytics_seattle_2 to remote-ip 5.6.7.8 from 10.1.1.1 ping-check off
```

History

This command was first available in ExtremeXOS 22.4.

Redundant remote IP addresses capability was added in ExtremeXOS 30.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable mlag port peer id

```
enable mlag port port peer peer_name id identifier
```

Description

Binds a local port or LAG to an .

Syntax Description

<i>port</i>	Specifies a local member port of the MLAG group.
<i>peer_name</i>	Specifies an alpha numeric string identifying the MLAG peer.
<i>identifier</i>	Specifies a unique MLAG identifier value. The range is 1 to 65000.

Default

N/A.

Usage Guidelines

Use this command to bind a local port or LAG to an MLAG that is uniquely identified by the MLAG ID value. The MLAG ID can be any number from 1 to 65000.

The specified port number may be a single port or the master port of a load sharing group but may not be a load sharing member port. If it is, a message similar to the following is displayed:

```
ERROR: Port 2 is a member of a load share group. Use the load share master port (10) instead.
```

A port can be part of only one MLAG. If you try to add it to another MLAG, a message similar to the following is displayed:

```
ERROR: Port 2 is already part of an MLAG Id 101
```

Once the MLAG group binding is made, any change to load sharing on MLAG ports is disallowed.

The MLAG peer must exist or the command will fail.

Example

The following command binds the local member port 2 to the peer switch switch101 with an identifier of 101:

```
# enable mlag port 2 peer switch101 id 101
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, 5720 series switches.

enable mlag port reload-delay

```
enable mlag port reload-delay
```

Description

This command enables reload-delay on Multi-switch Link Aggregation Group (MLAG) ports.

Syntax Description

This command has no arguments or variables.

Default

MLAG reload-delay is disabled by default.

Usage Guidelines

There are cases where MLAG ports comes up quicker than ISC ports after a switch reboot causing traffic loss during this time gap. After using the [configure mlag ports reload-delay](#) on page 863 command to configure a time delay for MLAG ports that provides enough time for ISC ports/neighborship of other Layer 3 protocols to come up, you have to issue this command to enable the timer.

Example

The following example enables the MLAG reload-delay timer:

```
# enable mlag port reload-delay
```

History

This command was first available in ExtremeXOS 22.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, 5720 series switches.

enable mld

```
enable mld {vlan vlan_name {MLDv1 | MLDv2} }
```

Description

Enables MLD on a router interface. If no VLAN is specified, MLD is enabled on all router interfaces.

Syntax Description

<i>vlan_name</i>	Specifies a VLAN name.
MLDv1	Sets the compatibility mode to MLDv1.
MLDv2	Sets the compatibility mode to MLDv2.

Default

Disabled.

Usage Guidelines

MLD is a protocol used by an IPv6 host to register its IPv6 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, IPv6 hosts respond to the query, and group registration is maintained.

MLD is disabled by default on the switch. However, the switch can be configured to enable the generation and processing of MLD packets. If compatibility mode is not specified in the command, MLDv1 compatibility mode is set.

A VLAN must have an IPv6 address to support MLD.

Example

The following example enables MLDv1 on the VLAN accounting:

```
enable mld vlan accounting
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 multicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

enable mld snooping

```
enable mld snooping {{vlan} vlan_name}
```

Description

Enables MLD snooping on the switch.

Syntax Description

<i>vlan_name</i>	Specifies a <u>VLAN</u> .
------------------	---------------------------

Default

Disabled.

Usage Guidelines

If a VLAN is specified, MLD snooping is enabled only on that VLAN, otherwise MLD snooping is enabled on all VLANs.

A VLAN must have an IPv6 address to support MLD.

Example

The following command enables MLD snooping on the switch:

```
enable mld snooping
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 multicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

enable mld snooping with-proxy

```
enable mld snooping with-proxy
```

Description

Enables the MLD snooping proxy.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

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 MLD leave messages so that they are forwarded only when the last member leaves the group.

This command can be used for troubleshooting purpose. It should be enabled for normal network operation. The command does not alter the snooping setting.

Example

The following command enables the MLD snooping proxy:

```
enable mld snooping with-proxy
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 multicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

enable mld ssm-map

```
enable mld ssm-map {{vr} vr_name}
```

Description

Enables MLD SSM mapping on a virtual router (VR).

Syntax Description

vr <i>vr_name</i>	Specifies a virtual router name.
--------------------------	----------------------------------

Default

Disabled.

Usage Guidelines

Use this command to enable MLD SSM mapping on a VR.

Configure the SSM address range using the `configure pim ipv6 ssm range [default | {policy} policy_name]` command before you enable SSM Mapping.

Example

The following example enables SSM mapping on VR1:

```
enable mld ssm-map vr vr1
```

History

This command was first available in ExtremeXOS 15.5.

Platform Availability

This command is available on the platforms listed for the IPv6 multicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

enable mpls

```
enable mpls
```

Description

Enables MPLS on the switch.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

Enabling MPLS allows MPLS processing to begin for any enabled MPLS protocols (RSVP-TE and/or LDP).

While MPLS is transitioning to the enabled state, only the MPLS show commands are accepted.

Before you can enable MPLS on a SummitStack, the stack must meet the following requirements:

- Each stack switch must meet the software and hardware requirements listed in the [Switch Engine 32.3 Feature License Requirements](#) document.
- You must configure the enhanced stacking protocol on each ExtremeSwitching series switch.



Note

When MPLS is enabled on a stack, you can only add MPLS-compatible switches to the stack.

Example

The following command globally enables MPLS on the switch:

```
enable mpls
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

enable mpls bfd

```
enable mpls bfd [{vlan} vlan_name | vlan all]
```

Description

Enables the Bidirectional Forwarding Detection (BFD) client for *MPLS* on the specified *VLAN* or all VLANs.

Syntax Description

<i>vlan_name</i>	Specifies the VLAN on which to enable the MPLS BFD client.
vlan all	Enables the MPLS BFD client on all VLANs.

Default

Disabled.

Usage Guidelines

This command causes MPLS to request a BFD session to each next-hop peer reachable through the named interface. BFD sessions are triggered by the establishment of an LSP over the interface. If this command is issued after LSPs are established, then the list of active LSPs is searched for next-hop peers associated with the named interface, and a BFD session is requested for each neighbor which does not already have a session. This command also instructs MPLS to begin to consider BFD neighbor session state updates as part of the effective interface link state reported to the MPLS upper layer protocols.



Note

BFD must be enabled on the interface before sessions can be established. To enable BFD, use the command: `[enable | disable] bfd vlan vlan_name` .

Example

The following command enables the MPLS BFD client on VLAN vlan1:

```
enable mpls bfd vlan1
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

enable mpls exp examination

```
enable mpls exp examination
```

Description

Enables assigning an *MPLS* packet to a *QoS* profile based on the MPLS packet's EXP value.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command enables assigning an MPLS packet to a QoS profile based on the MPLS packet's EXP value. The EXP values to QoS profile mappings are configured using the `configure mpls exp examination` command.

When disabled, all received MPLS packets are assigned to QoS profile qp1.

Example

The following command enables assignment of an MPLS packet to a QoS profile based on the MPLS packet's EXP value:

```
enable mpls exp examination
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

enable mpls exp replacement

```
enable mpls exp replacement
```

Description

Enables setting an MPLS packet's EXP value based on the packet's QoS profile.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command enables setting an MPLS packet's EXP value based on the packet's QoS profile. The QoS profiles to EXP value mappings are configured using the `configure mpls exp replacement` command.

When disabled, all MPLS packets are transmitted with an EXP value of zero.

Example

The following command enables the setting of an MPLS packet's EXP value based on the packet's QoS profile:

```
enable mpls exp replacement
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

enable mpls ldp bgp-routes

```
enable mpls ldp bgp-routes
```

Description

Enables LDP to use IP prefixes learned from *BGP* when establishing LDP LSPs.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

This command allows LDP to use routes learned via BGP when establishing LDP LSPs. Because each established LSP consumes internal resources, it is recommended that this setting be used only in BGP environments where the number of BGP routes is controlled.

When disabled, LDP does not establish LSPs to routes learned via BGP, thus reducing the internal resources used by LDP. Note that *MPLS* LSPs can still be used to transport packets to routes learned via BGP through the use of the `enable bgp mpls-next-hop` command.

Example

The following command enables the use of BGP routes by LDP:

```
enable mpls ldp bgp-routes
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

enable mpls ldp loop-detection

```
enable mpls ldp loop-detection
```

Description

Enables LDP loop detection on the switch.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

Loop detection provides a mechanism for finding looping LSPs and for preventing Label Request messages from looping in the presence of non-merge capable LSRs. The mechanism makes use of Path Vector and Hop Count TLVs carried by Label Request and Label Mapping messages.

Example

The following command globally enables LDP loop detection on the switch:

```
enable mpls ldp loop-detection
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

enable mpls ldp

```
enable mpls ldp [{vlan} vlan_name | vlan all]
```

Description

Enables LDP for the specified [MPLS](#) configured [VLANs](#).

Syntax Description

vlan	Enables LDP for one or more specific VLANs.
<i>vlan_name</i>	Enables LDP on the specified VLAN.
vlan all	Enables LDP for all VLANs that have been added to MPLS.

Default

Disabled.

Usage Guidelines

When LDP is enabled, LDP attempts to establish peer sessions with neighboring routers on the enabled VLAN. Once a peer session is established, LDP advertises labels for local IP interfaces and for routes learned from other neighboring routers. By default, LDP is disabled for all VLANs that have been added to MPLS. Specifying the optional **all** keyword enables LDP for all MPLS configured VLANs.

Example

The following command enables LDP for all VLANs that have been added to MPLS:

```
enable mpls ldp vlan all
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

enable mpls php

When enabled, PHP is requested on all LSPs advertised over that *VLAN* for which the switch is the egress LSR.

```
enable mpls php [{vlan} vlan_name | vlan all]
```

Description

Enables penultimate hop popping (PHP) on the specified VLAN.

Syntax Description

vlan	Enables PHP for one or more specific VLANs.
<i>vlan_name</i>	Enables PHP on the specified VLAN.
vlan all	Enables PHP for all VLANs that have been added to <i>MPLS</i> .

Default

Disabled.

Usage Guidelines

Penultimate hop popping is requested by assigning the Implicit Null Label in an advertised mapping. Extreme's MPLS implementation always performs penultimate hop popping when requested to do so by a peer LSR. When the all VLANs option is selected, PHP is enabled on all configured VLANs that have been added to MPLS.

Example

The following command enables penultimate hop popping (PHP) on the specified VLAN:

```
enable mpls php vlan vlan1
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

enable mpls protocol ldp

```
enable mpls protocol ldp
```

Description

Enables LDP for the switch.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

When LDP is enabled, LDP attempts to establish peer sessions with neighboring routers on [VLAN](#) interfaces where LDP has been enabled . Once a peer session is established, LDP can advertise labels for local IP interfaces and for routes learned from other neighboring routers. While LDP is transitioning to the enabled state, only the [MPLS](#) show commands are accepted.

Note that the LDP protocol must be enabled to establish VPLS pseudo-wires even if the transport LSPs are being established using RSVP-TE.

Example

The following command globally enables LDP on the switch:

```
enable mpls protocol ldp
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

enable mpls protocol rsvp-te

```
enable mpls protocol rsvp-te
```

Description

Enables RSVP-TE for the switch.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

When RSVP-TE is enabled, configured LSPs begin the process of TE LSP establishment and [VLAN](#) interfaces that have RSVP-TE enabled begin processing RSVP path/reserve messages. By default, RSVP-TE is disabled. While RSVP-TE is transitioning to the enabled state, only the [MPLS](#) show commands are accepted.

Example

The following command globally enables RSVP-TE on the switch:

```
enable mpls protocol rsvp-te
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

enable mpls rsvp-te bundle-message

```
enable mpls rsvp-te bundle-message [{vlan} vlan_name | vlan all]
```

Description

Enables the bundling of RSVP-TE messages for the specified [VLAN](#) interface.

Syntax Description

vlan	Specifies that message-bundling is to be enabled on one or more VLAN interfaces.
<i>vlan_name</i>	Identifies a VLAN interface for which message bundling is to be enabled.
vlan all	Specifies that message bundling is to be enabled on all VLANs that have been added to MPLS .

Default

Disabled.

Usage Guidelines

Enabling message bundling can improve control plane scalability by allowing the switch to bundle multiple RSVP-TE messages into a single PDU. Not all devices support bundled messages. If the switch determines that a peer LSR, connected to a specific interface, does not support message bundling, the switch reverts to sending separate PDUs for each message on that interface. By default, message bundling is disabled. Specifying the **all** keyword enables message bundling on all MPLS-configured VLANs.

Example

The following command enables message bundling on the specified VLAN:

```
enable mpls rsvp-te bundle-message vlan vlan_1
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

enable mpls rsvp-te fast-reroute

```
enable mpls rsvp-te fast-reroute
```

Description

Enables the MPLS RSVP-TE fast reroute (FRR) protection feature.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

You can configure FRR LSPs only when FRR is enabled on the LSR. Enabling FRR protection on the LSR automatically enables the point-of-local-repair and merge-point capabilities on the LSR. FRR should be enabled on all LSRs along each FRR LSP path.

Example

The following command enables FRR protection on the local switch:

```
enable mpls rsvp-te fast-reroute
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

enable mpls rsvp-te lsp

```
enable mpls rsvp-te lsp [lsp_name | all]
```

Description

Enables one or more RSVP-TE LSPs.

Syntax Description

<i>lsp_name</i>	Specifies the ingress LSP within the switch to be enabled.
all	Enables all RSVP-TE configured ingress LSPs.

Default

Enabled.

Usage Guidelines

When an RSVP-TE LSP is enabled, the switch attempts to set up the LSP by signaling the destination by sending a path message using the assigned path and profile. By default, all newly created LSPs are enabled and can become active when the LSP has been configured. Note that an LSP must be configured with at least one path before it can be signaled.

Example

The following command enables all RSVP-TE-configured LSPs:

```
enable mpls rsvp-te lsp all
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

enable mpls rsvp-te summary-refresh

```
enable mpls rsvp-te summary-refresh [{vlan} vlan_name | vlan all]
```

Description

Enables the sending of summary refresh messages, instead of path messages, to refresh RSVP-TE path state for the specified [VLAN](#) interface.

Syntax Description

vlan	Specifies that summary refresh messages are to refresh the RSVP-TE path state on one or more VLAN interfaces.
<i>vlan_name</i>	Identifies a VLAN interface on which RSVP-TE summary refresh messages are to refresh the RSVP-TE path state.
vlan all	Specifies that summary refresh messages are to refresh the RSVP-TE path state on all VLANs that have been added to MPLS .

Default

Disabled.

Usage Guidelines

Enabling summary refresh can improve control plane scalability by refreshing multiple LSPs in a single message. Not all devices support summary refresh. If the switch determines that a peer LSR, connected to a specific interface, does not support summary refresh, the switch reverts to using path messages on that interface. By default, summary refresh is disabled. Specifying the **all** keyword enables summary refresh on all MPLS-configured VLANs.

Example

The following command enables summary refresh on the specified VLAN:

```
enable mpls rsvp-te summary-refresh vlan vlan_1
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

enable mpls rsvp-te

```
enable mpls rsvp-te [{vlan} vlan_name | vlan all]
```

Description

Enables RSVP-TE for the specified *MPLS*-configured *VLAN*.

Syntax Description

vlan	Specifies that RSVP-TE is to be enabled on one or more VLANs.
<i>vlan_name</i>	Identifies a specific VLAN on which RSVP-TE is to be enabled.
vlan all	Enables RSVP-TE on all VLANs that have been added to MPLS.

Default

Disabled.

Usage Guidelines

When RSVP-TE is enabled, TE LSP establishment for configured LSPs begins and the processing of RSVP path/reserve messages from peer LSRs is permitted. By default, RSVP-TE is disabled for all MPLS-configured VLANs. Specifying the optional all keyword enables RSVP-TE for all VLANs that have been added to MPLS.

Example

The following command enables RSVP-TE on the specified VLAN:

```
enable mpls rsvp-te vlan vlan_1
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

enable mpls static lsp

```
enable mpls static lsp {lsp_name | all }
```

Description

Administratively enables one or all static LSPs.

Syntax Description

<i>lsp_name</i>	Identifies the LSP to be enabled.
all	Specifies that all static LSPs on this LSR are to be enabled.

Default

N/A.

Usage Guidelines

On executing this command, the software tries to activate the static LSP by programming the LSP in hardware. Static LSPs are not enabled by default. You need to explicitly enable LSPs after the ingress and egress segments have been configured.

Example

The following command enables a static LSP:

```
enable mpls static lsp lsp598
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in the [Switch Engine 32.3 Feature License Requirements](#) document.[ture-link-in-22.1"/>](#)

enable mpls vlan

```
enable mpls [{vlan} vlan_name | vlan all]
```

Description

Enables the [MPLS](#) interface for the specified [VLAN](#).

Syntax Description

vlan	Enables an MPLS interface for one or more specific VLANs.
<i>vlan_name</i>	Enables an MPLS interface on the specified VLAN.
vlan all	Enables an MPLS interface for all VLANs that have been added to MPLS.

Default

Disabled.

Example

The following command enables an MPLS interface for the specified VLAN:

```
enable mpls vlan vlan-nyc
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

enable msdp data-encapsulation

```
enable msdp data-encapsulation {vr vrname}
```

Description

Enables the encapsulation of locally originated SA messages with multicast data (if available).

Syntax Description

<i>vrname</i>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.
---------------	---

Default

By default, multicast data packet encapsulation is enabled for locally originated SA messages. Multicast data packets with a packet size of up to 8 KB are encapsulated in SA messages.

Usage Guidelines

Enable data encapsulation to handle bursty sources.

Example

The following command enables multicast data packet encapsulation:

```
enable msdp data-encapsulation
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the *MSDP* feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

enable msdp export local-sa

```
enable msdp export local-sa {export-filter filter-name} {vr vrname}
```

Description

Enables the advertisement of local sources to groups for which the router is an RP.

Syntax Description

<i>filter-name</i>	Specifies the policy to associate with the export of local sources. No policy is specified by default.
<i>vrname</i>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

By default, the export of local sources is enabled. All sources are advertised if the router is an RP for the groups.

Usage Guidelines

You can create a policy to filter out some of the local sources so that they are not advertised to *MSDP* peers and exposed to the external multicast domain. To configure an export filter, you must first disable the export of local sources (with the `disable msdp export local-sa` command), and then re-enable it with an export filter (with the `enable msdp export local-sa export-filter` command).

You can use the following policy attributes in an export policy. All other attributes are ignored.

- Match:
 - multicast-group
 - multicast-source

- pim-rp
- Set:
 - permit
 - deny

Please note that the syntax for “multicast-group”, “multicast-source,” and “pim-rp” are the same as for the “nlri” policy attribute.

```
[multicast-group | multicast-source | pim-rp] [ipaddress | any]/mask-length> {exact}
[multicast-group | multicast-source | pim-rp] [ipaddress | any] mask mask {exact}
```

An example of an MSDP policy file follows:

```
entry allow_internal_rp {
  if match any {
    multicast-group 234.67.89.0/24;
    multicast-source 23.123.45.0/24;
    pim-rp 10.203.134.5/32;
  } then {
    permit;
  }
}
entry deny_local_group239 {
  if match any {
    multicast-group 239.0.0.0/8;
    multicast-source 23.123.45.0/24;
  } then {
    deny;
  }
}
entry allow_external_rp_172 {
  if {
    multicast-group 234.172.0.0/16;
  } then {
    permit
  }
}
# deny remaining entries
```

Example

The following command enables the advertisement of local sources:

```
nable msdp export local-sa
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

enable msdp peer

```
enable msdp [{peer} remoteaddr | peer all] {vr vr_name}
```

Description

Configures the administrative state of an *MSDP* peer.

Syntax Description

all	Enables all MSDP peers.
<i>remoteaddr</i>	Specifies the IP address of the MSDP peer to configure.
<i>vr_name</i>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

By default, MSDP peers are disabled.

Usage Guidelines

You must use this command to administratively enable the MSDP peers before they can establish peering sessions and start exchanging SA messages.

Example

The following example enables an MSDP peer:

```
enable msdp peer 192.168.45.43
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

enable msdp process-sa-request

```
enable msdp [{peer} remoteaddr | peer all] process-sa-request {sa-request-filter filter-name } {vr vr_name}
```

Description

This command configures *MSDP* to receive and process SA request messages from a specified peer or all peers. If an SA request filter is specified, only SA request messages from those groups permitted are accepted. All others are ignored.

Syntax Description

peer all	Specifies all MSDP peers.
<i>filter-name</i>	Specifies the name of the policy filter associated with SA request processing.
<i>remoteaddr</i>	Specifies the IP address of the MSDP peer.
<i>vr_name</i>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

By default, all SA request messages are accepted from peers.

Usage Guidelines

Use this command to configure the router to accept all or just some SA request messages from peers. If no policy is specified, all SA request messages are accepted. If a policy is specified, only SA request messages from those groups permitted are accepted, and all others are ignored.

You cannot change an SA request filter while SA request processing is enabled for an MSDP peer. You must first disable SA request processing for a peer and then re-enable it with an SA request filter.

You can use the following policy attributes in an SA request policy. All other attributes are ignored.

- Match:
 - multicast-group
 - multicast-source
 - pim-rp
- Set:
 - permit
 - deny

Example

The following example enables processing of SA request messages received from a peer with the IP address 192.168.45.43:

```
enable msdp peer 192.168.45.43 process-sa-request sa-request-filter intra_domain
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

enable msdp

```
enable msdp {vr vrname}
```

Description

Enables *MSDP* on a virtual router.

Syntax Description

<i>vrname</i>	Specifies the name of the virtual router on which MSDP is being enabled or disabled. If a name is not specified, it is extracted from the current CLI context.
---------------	--

Default

MSDP is disabled by default.

Usage Guidelines

None.

Example

The following command enables MSDP on a virtual router:

```
enable msdp
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

enable msrp ports

```
enable msrp ports [port_list | all]
```

Description

Enables MSRP in the ports listed in the command after the keyword **ports**.

Syntax Description

msrp	Multiple Stream Registration Protocol.
<i>port_list</i>	Port list separated by a comma or "-".
all	All ports.

Default

Disabled.

Usage Guidelines

Use this command to enable MSRP in the ports listed or all ports.



Note

MSRP is not supported for Link Aggregated Ports.

Example

```
# enable msrp ports 1-3
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on all platforms that support the AVB feature and that have an AVB feature pack license installed. To see which platforms support AVB and for information about obtaining a license, see the [Switch Engine 32.3 Feature License Requirements](#).

enable msrp

```
enable msrp
```

Description

Enables MSRP globally on the switch.

Syntax Description

msrp	Multiple Stream Registration Protocol.
-------------	--

Default

Disabled.

Usage Guidelines

Use this command to enable MSRP globally on a switch.

Example

The following command enables MSRP:

```
enable msrp
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on all platforms that support the AVB feature and that have an AVB feature pack license installed. To see which platforms support AVB and for information about obtaining a license, see the [Switch Engine 32.3 Feature License Requirements](#).

enable mvr

```
enable mvr
```

Description

Enables MVR on the system.

Syntax Descripton

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

None.

Example

The following command enables MVR on the system:

```
enable mvr
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MVRP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

enable mvrp

```
enable mvrp
```

Description

Enables MVRP globally on a switch.

Syntax Description

mvrp	Multiple <u>VLAN</u> Registration Protocol.
-------------	---

Default

Disabled.

Usage Guidelines

Use this command to enable MVRP globally on a switch. MVRP is run on the MVRP enabled ports only if the global setting is enabled. By default, MVRP is disabled globally and on individual ports. When MVRP is disabled globally, all MVRP packets are forwarded transparently.

Example

The following command enables MVRP globally on the switch:

```
enable mvrp
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable mvrp ports

```
enable mvrp ports [port_list | all]
```

Description

Enables MVRP on a given set of ports.

Syntax Description

mvrp	Multiple <u>VLAN</u> Registration Protocol
<i>port_list</i>	Port(s) on which MVRP is to be enabled.
all	All ports.

Default

Disabled.

Usage Guidelines

Use this command to enable MVRP on given set of ports. MVRP is run on the MVRP enabled ports only if the global setting is enabled. By default, MVRP is disabled globally and on individual ports. When MVRP is disabled globally, all MVRP packets will be forwarded transparently. An error message is displayed if the user tries to enable/disable MVRP on a lag member port which is not the master port. No configuration changes are made.

Example

The following command enables MVRP on ports 4 and 5:

```
enable mvrp ports 4-5
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable neighbor-discovery refresh

```
enable neighbor-discovery {vr vr_name} refresh
```

Description

Enables the IPv6 neighbor cache to refresh each entry before the timeout period expires.

Syntax Description

<i>vr_name</i>	Specifies a VR or VRF.
----------------	------------------------

Default

Enabled.

Usage Guidelines

None.

Example

The following example enables the refresh of neighbor discovery cache entries:

```
enable neighbor-discovery refresh
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

enable netlogin

```
enable netlogin [{dot1x} {mac} {web-based}]
```

Description

Enables network login authentication modes.

Syntax Description

dot1x	Specifies 802.1X authentication.
mac	Specifies MAC-based authentication.
web-based	Specifies web-based authentication.

Default

All types of authentication are disabled.

Usage Guidelines

Any combination of types of authentication can be enabled on the same switch. At least one of the authentication types must be specified on the command line.

Entering **enable netlogin mac** adds **configure netlogin add mac-list default** configuration by default.

To disable an authentication mode, use the following command:

```
disable netlogin [{dot1x} {mac} {web-based}]
```

Example

The following command enables web-based network login:

```
enable netlogin web-based
```

History

This command was first available in ExtremeXOS 11.1.

Default **configure netlogin add mac-list default** configuration was added in ExtremeXOS 31.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable netlogin authentication failure vlan ports

```
enable netlogin authentication failure vlan ports [ports | all]
```

Description

Enables the configured authentication failure VLAN on the specified ports.

Syntax Description

all	Specifies all ports included in the authentication failure VLAN.
<i>ports</i>	Specifies one or more ports or slots and ports on which the authentication failure VLAN is enabled.

Default

All ports.

Usage Guidelines

Use this command to enable the configured authentication failure VLAN on either the specified ports, or all ports.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable netlogin authentication service-unavailable vlan ports

```
enable netlogin authentication service-unavailable vlan ports [ports | all]
```

Description

Enables the configured authentication service-unavailable VLAN on the specified ports.

Syntax Description

ports	Specifies one or more ports or slots and ports on which the service-unavailable VLAN is enabled.
all	Specifies all ports included in the service-unavailable VLAN.

Default

All ports.

Usage Guidelines

Use this command to enable the configured authentication service-unavailable VLAN on the specified ports, or on all ports.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable netlogin dot1x guest-vlan ports

```
enable netlogin dot1x guest-vlan ports [all | ports]
```

Description

Enables the guest VLAN on the specified 802.1X network login ports.

Syntax Description

all	Specifies all ports included in the guest VLAN.
<i>ports</i>	Specifies one or more ports or slots and ports on which the guest VLAN is enabled.

Default

Disabled.

Usage Guidelines

A guest VLAN provides limited or restricted network access if a supplicant connected to a port does not respond to the 802.1X authentication requests from the switch. A port always moves untagged into the guest VLAN.

Modifying the Supplicant Timer

By default, the switch attempts to authenticate the supplicant every 30 seconds for a maximum of three tries. If the supplicant does not respond to the authentication requests, the client moves to the guest VLAN. The number of authentication attempts is a user-configured parameter with allowed values in the range of 1 to 10.

To modify the supplicant response timer, use the following command and specify the supp-resp-timeout parameter:

```
configure netlogin dot1x timers [{server-timeout server_timeout}  
{quiet-periodquiet_period} {reauth-periodreauth_period} {reauth-  
maxmax_num_reauths}} {supp-resp-timeoutsupp_resp_timeout}
```

Creating the Guest VLAN

Before you can enable the guest VLAN on the specified ports, you must create the guest VLAN. To create the guest VLAN, use the following command:

```
configure netlogin dot1x guest-vlan vlan_name {portsport_list}
```

Example

The following command enables the guest VLAN on all ports:

```
enable netlogin dot1x guest-vlan ports all
```

The following command enables the guest VLAN on ports 2 and 3:

```
enable netlogin dot1x guest-vlan ports 2,3
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable netlogin logout-privilege

```
enable netlogin logout-privilege
```

Description

Enables network login logout pop-up window.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

This command controls the logout window pop-up on the web-based network client. This command applies only to the web-based authentication mode of network login.

Example

The following command enables network login logout-privilege:

```
enable netlogin logout-privilege
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable netlogin ports

```
enable netlogin ports ports [{dot1x} {mac} {web-based}]
```

Description

Enables NetLogin on a specified port for a particular authentication method.

Syntax Description

<i>ports</i>	Specifies the ports for which NetLogin should be enabled.
dot1x	Specifies 802.1X authentication.
mac	Specifies MAC-based authentication.
web-based	Specifies web-based authentication.

Default

All methods are disabled on all ports.

Usage Guidelines

For campus mode NetLogin with web-based clients, 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 NetLogin.
- The switch must be configured as a [*RADIUS*](#) client, and the RADIUS server must be configured to enable the NetLogin capability.

For ISP mode login, no special conditions are required. A RADIUS server must be used for authentication.

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

Windows authentication is not supported via NetLogin.

To support NetLogin on all user virtual routers (VRs) in policy mode, remove any associated VRs from the port before enabling NetLogin (see [configure vr delete ports](#) on page 1564). This is applicable for uplink ports and ISC ports. This must be done prior to authentication so that once the client gets authenticated the ports can move across different VLANs of various VRs.

Example

The following command configures NetLogin on port 2:9 using web-based authentication:

```
enable netlogin ports 2:9 web-based
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable netlogin reauthentication-on-refresh

```
enable netlogin reauthentication-on-refresh
```

Description

Enables network login reauthentication on refresh.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

The web-based Netlogin client's session is periodically refreshed by sending a HTTP request which acts as a keep-alive without actually re-authenticating the user's credentials with the back-end RADIUS server or local database. If reauthenticate-on-refresh is enabled, re-authentication occurs with the session refresh.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable netlogin redirect-page

```
enable netlogin redirect-page
```

Description

Enables the network login redirect page function.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

This command enables the network login redirect page so that the client is sent to the redirect page rather than the original page.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable netlogin session-refresh

```
enable netlogin session-refresh {refresh_minutes}
```

Description

Enables network login session refresh.

Syntax Description

<i>refresh_minutes</i>	Specifies the session refresh time for network login in minutes.
------------------------	--

Default

Enabled, with a value of three minutes for session refresh.

Usage Guidelines

Network login sessions can refresh themselves after a configured timeout. After the user has been logged in successfully, a logout window opens which can be used to close the connection by clicking on the Logout link. Any abnormal closing of this window is detected on the switch and the user is logged out after a time interval as configured for session refresh. The session refresh is enabled and set to three minutes by default. The value can range from 1 to 255 minutes. When you configure the network login session refresh for the logout window, ensure that the *FDB* aging timer is greater than the network login session refresh timer.

This command applies only to the web-based authentication mode of network login.

To reset the session refresh value to the default behavior, use this command without the minutes parameter.

Example

The following command enables network login session refresh and sets the refresh time to ten minutes:

```
enable netlogin session-refresh 10
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable network-clock gptp

```
enable network-clock gptp
```

Description

Enables gPTP on the switch.

Syntax Description

network-clock	Network clock.
gptp	IEEE 802.1AS Generalized Precision Time Protocol (gPTP).

Default

Disabled.

Usage Guidelines

Use this command to enable gPTP.

Example

```
# enable network-clock gptp
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on all platforms that support the AVB feature and that have an AVB feature pack license installed. To see which platforms support AVB and for information about obtaining a license, see the [Switch Engine 32.3 Feature License Requirements](#).

enable network-clock gtp ports

```
enable network-clock gtp ports [port_list {only} | all]
```

Description

Enables gPTP on one or more ports.

Syntax Description

<i>port_list</i>	Specifies one or more of the switch's physical ports.
only	Apply change only to specified port, even if port is master of a load sharing group.
all	Specifies all of the switch's physical ports.

Default

Disabled.

Usage Guidelines

Use this command to configure on which ports gPTP runs. gPTP does not run on any ports if it is not first enabled in the switch by the `enable network-clock gtp` command.

Example

```
# enable network-clock gtp ports 4
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on the ExtremeSwitching 5520 switch if the AVB feature pack license is installed on the switch..

enable nodealias ports

```
enable nodealias ports [port_list | all]
```

Description

This command enables the Node Alias feature on specified ports. Node Alias discovers information about the end systems on a per-port basis. Information from packets from end systems, such as VLANID, source MAC address, source IP address, protocol, etc. are captured in a database that can be queried.

Syntax Description

nodealias	Node Alias feature that maps source IP address, MAC address, host name, and protocol on a per port basis.
ports	Designates that Node Alias should be enabled on specified ports.
<i>port_list</i>	Specifies on which ports to have Node Alias enabled. Designated as a port list separated by comma (,) or dash (-).
all	Specifies that all ports have Node Alias enabled.

Default

Node Alias is disabled by default on all ports.

Usage Guidelines

If the port is part of a [LAG](#), Node Alias should be enabled separately on each LAG port.

Example

The following example enables Node Alias on all ports:

```
enable nodealias ports all
```

History

This command was first available in ExtremeXOS 22.2.

Platform Availability

VLAN

enable nodealias protocol

```
enable nodealias protocol [protocol_name | all]
```

Description

This command designates the specific protocols detected for the Node Alias feature. Node Alias discovers information about the end systems on a per-port basis. Information from packets from end systems, such as VLANID, source MAC address, source IP address, protocol, etc. are captured in a database that can be queried.

Syntax Description

nodealias	Node Alias feature that maps source IP address, MAC address, host name, and protocol on a per port basis.
protocol	Designates selection of protocols to detect.
<i>protocol_name</i>	Specifies enabling a protocol to detect (one at a time). The following protocols are enabled by default: IPv4, IPv6, <i>OSPF</i> , <i>BGP</i> , <i>VRRP</i> , DHCP, DHCP, BOOTP, BOOTP, UDP, BPDU, LLMNR, SSDP, and mDNS.
any	Specifies enabling all Node Alias-supported protocols.

Default

The following protocols are enabled by default: IPv4, IPv6, OSPF, BGP, VRRP, DHCP, DHCP, BOOTP, BOOTP, UDP, BPDU, LLMNR, SSDP, and mDNS.



Note

- ARP is categorized under IP.
- UDP entry is created when destination IP address is broadcast.
- BPDU means *STP* and GVRP frames.

Usage Guidelines

By default, the following protocols are enabled (IPv4, IPv6, OSPF, BGP, VRRP, DHCP, DHCP, BOOTP, BOOTP, UDP, BPDU, LLMNR, SSDP, mDNS). You can optionally disable any of these protocols (and then enable them back if desired).

Example

The following example specifically enables BGP to be detected:

```
enable nodealias protocol bgp
```

History

This command was first available in ExtremeXOS 22.2.

Platform Availability

VLAN

enable ntp

```
enable ntp
```

Description

Enables NTP globally on the switch.

Syntax Description

N/A.

Default

NTP is disabled by default.

Usage Guidelines

N/A.

Example

The following command enables NTP globally on the switch:

```
enable ntp
```

History

This command was first available in ExtremeXOS 12.7.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable ntp authentication

```
enable ntp authentication
```

Description

Enables NTP authentication globally on the switch.

Syntax Description

N/A.

Default

NTP authentication is disabled by default.

Usage Guidelines

If authentication is disabled, NTP will not use any authentication mechanism to a server or from clients. To use authentication for a specific server, enable NTP authentication globally, and then configure an RSA Data Security, Inc. [MD5](#) Message-Digest Algorithm or SHA256 key index for the specific server.

Example

The following command enables NTP authentication globally on the switch:

```
enable ntp authentication
```

History

This command was first available in ExtremeXOS 12.7.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable ntp broadcast-client

```
enable ntp broadcast-client [{vr} vr_name]
```

Description

Enables an NTP broadcast client on the switch.

Syntax Description

broadcast-client	Specifies enabling NTP broadcast client.
vr	Specifies enabling NTP broadcast client for a VR.
<i>vr_name</i>	Specifies the VR name. If a VR name is not specified, the VR of current command context is used.

Default

An NTP broadcast client is enabled by default.

If a VR name is not specified, the VR of current command context is used.

Usage Guidelines

If the broadcast client function is enabled, the system can receive broadcast-based NTP messages and process them only if a [VLAN](#) is enabled for NTP and the VLAN is active.

Example

The following command enables an NTP broadcast client on the switch:

```
enable ntp broadcast-client
```

History

This command was first available in ExtremeXOS 12.7.

The **vr** was added in ExtremeXOS 22.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable ntp broadcast-server

```
enable ntp {vlan} vlan-name broadcast-server {key keyid}
```

Description

Enables NTP to send broadcast messages with or without a key to a VLAN.

Syntax Description

<i>vlan-name</i>	Specifies the name of a particular VLAN on which to enable or disable NTP.
<i>keyid</i>	Specifies the key ID as a value from 1 to 65534.

Default

An NTP broadcast server is enabled by default.

Usage Guidelines

For the broadcast server function to work correctly, configure a VLAN to forward broadcast packets by using the `enable ipforwarding broadcast vlan-name` command. All broadcast clients will receive clock information from the broadcasted clock messages.

Example

The following command enables an NTP broadcast server on the switch:

```
enable ntp vlan toSW3 broadcast-server key 100
```

History

This command was first available in ExtremeXOS 12.7.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable ntp vlan

```
enable ntp [{vlan} vlan-name | all] [{vr} vr_name]
```

Description

Enables NTP on a VLAN.

Syntax Description

enable	Enables NTP on a VLAN.
<i>vlan-name</i>	Specifies the name of a particular VLAN on which to enable or disable NTP.
all	Enables or disables NTP on all VLANs.
vr	Specifies setting up NTP on a VR.
<i>vr_name</i>	Specifies the VR name to enable NTP on. If a VR name is not specified, the VR of current command context is used.

Default

NTP is disabled on all VLANs by default.

Usage Guidelines

N/A.

Example

The following command enables NTP on a VLAN named "Southwest":

```
enable ntp vlan Southwest
```

History

This command was first available in ExtremeXOS 12.7.

The **vr** option was added in ExtremeXOS 22.2

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable ntp vr

```
enable ntp vr vr_name
```

Description

This command enables and configures NTP for the specified VR.

Syntax Description

vr	Specifies setting up NTP on a VR.
<i>vr_name</i>	Specifies the VR name to enable NTP on. If a VR name is not specified, the VR of current command context is used.

Default

If a VR name is not specified, the VR of current command context is used.

Example

The following example enables NTP on a VR named "vr1".

```
enable ntp vr vr1
```

History

This command was first available in ExtremeXOS 22.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable ospf

```
enable ospf
```

Description

Enables the *OSPF* process for the router.

Syntax Description

This command has no keywords or arguments.

Default

N/A.

Usage Guidelines

Not applicable.

Example

The following command enables the OSPF process for the router:

```
enable ospf
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [***Switch Engine 32.3 Feature License Requirements***](#) document.

enable ospf capability opaque-lsa

```
enable ospf capability opaque-lsa
```

Description

Enables opaque LSAs across the entire system.

Syntax Description

This command has no keywords 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 first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [*Switch Engine 32.3 Feature License Requirements*](#) document.

enable ospf export

```
enable ospf export [bgp | direct | e-bgp | i-bgp | rip | static | isis
| isis-level-1 | isis-level-1-external | isis-level-2 | isis-level-2-
external | host-mobility] [cost cost type [ase-type-1 | ase-type-2]
{tag number} | policy-map]
```

Description

Enables redistribution of routes to [*OSPF*](#).

Syntax Description

bgp	Specifies <i>BGP</i> routes.
direct	Specifies direct routes.
e-bgp	Specifies E-BGP routes.
i-bgp	Specifies I-BGP routes.
rip	Specifies <i>RIP</i> routes.
static	Specifies static routes.

isis	Specifies IS-IS routes.
isis-level-1	Specifies ISIS Level 1 routes.
isis-level-1-external	Specifies ISIS Level 1 External routes.
isis-level-2	Specifies ISIS Level 2 routes.
isis-level-2-external	Specifies ISIS Level 2 External routes.
host-mobility	Specifies host-mobility routes.
<i>cost</i>	Specifies a cost metric.
ase-type-1	Specifies AS-external type 1 routes.
ase-type-2	Specifies AS-external type 2 routes.
<i>number</i>	Specifies a tag value.
<i>policy-map</i>	Specifies a policy.

Default

The default tag number is 0. The default setting is disabled.

Usage Guidelines

After OSPF export is enabled, the OSPF router is considered to be an ASBR. Interface routes that correspond to the interface that has OSPF enabled are ignored.

The cost metric is inserted for all BGP, IS-IS, 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, type, and tag values can be inserted for all the export routes, or a policy can be used for selective insertion. When a policy is associated with the export command, the policy is applied on every exported route. The exported routes can also be filtered using a policy.

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 ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

NEW! enable ospf export vr

```
enable ospf export {vr} vr-name route-type [policy-map | cost cost type
ase-type-1 | ase-type-2] [tag number}] [exclude-private]
```

Description

Enables redistribution of routes between *OSPF* instances.

Syntax Description

vr	Specifies Virtual Router.
<i>vr-name</i>	Specifies the source Virtual Router for exported routes.
<i>route-type</i>	Specifies the OSPF inter-VR export route type.
<i>policy-map</i>	Specifies the route map.
cost	Specifies cost.
<i>cost</i>	Specifies the cost metric. Range is 0-65535.
type	Specifies route type.
ase-type-1	Specifies AS-external type 1 routes.
ase-type-2	Specifies AS-external type 2 routes.
tag	Specifies a tag.
<i>number</i>	Specifies a tag value.
exclude-private	Specifies to exclude routes corresponding to private IP ranges. Default is include private.

Default

N/A

Usage Guidelines

This command is executed in the destination VR and specifies the source VR.

The **exclude-private** option excludes prefixes in the private IP address ranges. Only OSPF route types can be exported between instances.

Example

The following command enables OSPF ??:

```
enable ospf export vr cost 1 ase-type-1 tag 0
```

History

This command was first available in ExtremeXOS 32.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable ospf mpls-next-hop

```
enable ospf mpls-next-hop {vr vrf_name}
```

Description

Enables IP forwarding over calculated *MPLS* LSPs to subnets learned through *OSPF*.

Syntax Description

<i>vrf-name</i>	Specifies OSPF on a particular VRF.
-----------------	-------------------------------------

Default

Disabled.

Usage Guidelines

This command enables IP forwarding over calculated MPLS LSPs to subnets learned through OSPF. (Calculated refers to an LSP that only reaches part of the way to the destination). By default, IP forwarding over MPLS LSPs to subnets learned via OSPF is disabled.

In order to configure OSPF on a particular VRF, you must supply the optional **vr** *vr-name* CLI parameter.

Example

The following command enables OSPF's use of MPLS LSPs to reach OSPF routes:

```
enable ospf mpls-next-hop
```

History

This command was first available in ExtremeXOS 11.6.

The **vr** keyword and *vrf_name* variable were added in ExtremeXOS 15.3.

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

enable ospf originate-default

```
enable ospf originate-default {always} cost cost type [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.
<i>cost</i>	Specifies a cost metric.
ase-type-1	Specifies AS-external type 1 routes.
ase-type-2	Specifies AS-external type 2 routes.
<i>number</i>	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 ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

enable ospf restart-helper-lsa-check

```
enable ospf [vlan [all | vlan-name] | area area-identifier |virtual-link
router-identifier area-identifier] restart-helper-lsa-check
```

Description

Enables the restart helper router to terminate graceful OSPF restart when received LSAs would affect the restarting router.

Syntax Description

all	Specifies all <u>VLANs</u>
<i>vlan-name</i>	Specifies a VLAN name.
<i>router-identifier</i>	Specifies the router ID of the remote router of the virtual link.
<i>area-identifier</i>	Specifies an OSPF area.

Default

The default is enabled.

Usage Guidelines

This command configures the restart helper router to terminate graceful OSPF restart when received LSAs would affect the restarting router. This will occur when the restart-helper receives an LSA that will be flooded to the restarting router or when there is a changed LSA on the restarting router's retransmission list when graceful restart is initiated.

Example

The following command configures a router to terminate graceful OSPF restart for all routers in area 10.20.30.40 if it receives an LSA that would affect routing:

```
enable ospf area 10.20.30.40 restart-helper-lsa-check
```

History

This command was first available in ExtremeXOS 11.3.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the Switch Engine 32.3 Feature License Requirements document.

enable ospf use-ip-router-alert

```
enable ospf use-ip-router-alert
```

Description

Enables the generation of the *OSPF* router alert IP option.

Syntax Description

This command has no keywords or arguments.

Default

Disabled.

Usage Guidelines

Not applicable.

Example

The following command enables the OSPF router alert IP option:

```
enable ospf use-ip-router-alert
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the ***Switch Engine 32.3 Feature License Requirements*** document.

enable ospfv3

```
enable ospfv3
```

Description

Enables *OSPFv3* for the router.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

When OSPFv3 is enabled, it will start exchanging Hellos on all of its active interfaces. It will also start exporting routes into OSPFv3 routing domain from other protocols, if enabled.

When OSPFv3 is disabled, it will release all the run-time allocated resources like adjacencies, link state advertisements, run-time memory, etc.

OSPFv3 can be enabled successfully if and only if:

- At least one of the VLANs in the current virtual router has one IPv4 address configured

—OR—

- You explicitly configure the OSPFv3 router ID, a four-byte, dotted decimal number

Example

The following command enables OSPFv3 for the router:

```
enable ospfv3
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [*Switch Engine 32.3 Feature License Requirements*](#) document.

enable ospfv3 export

```
enable ospfv3 export [direct | ripng | static | isis | isis-level-1  
| isis-level-1-external | isis-level-2 | isis-level-2-external | bgp  
| e-bgp | i-bgp | host-mobility] [cost cost type [ase-type-1 | ase-  
type-2] | policy_map]
```

Description

Enables redistribution of routes to [*OSPFv3*](#).

Syntax Description

direct	Specifies direct routes.
ripng	Specifies <i>RIPng</i> routes.

static	Specifies static routes.
isis	Specifies IS-IS routes.
isis-level-1	Specifies IS-IS Level 1 routes.
isis-level-1-external	Specifies IS-IS Level 1 External routes.
isis-level-2	Specifies IS-IS Level 2 routes.
isis-level-2-external	Specifies IS-IS Level 2 External routes.
bgp	Specifies <i>BGP</i> IPv6 routes.
i-bgp	Specifies internal BGP IPv6 routes.
e-bgp	Specifies external BGP IPv6 routes.
host-mobility	Specifies host-mobility routes.
<i>cost</i>	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.
<i>policy_map</i>	Specifies a policy.

Default

The default setting is disabled.

Usage Guidelines

The cost metric is inserted for all RIPng-learned, static, and direct routes injected into OSPFv3. If the cost metric is set to 0, the cost is inserted from the route.

The same cost and type values can be inserted for all the export routes, or a policy can be used for selective insertion. When a policy is associated with the export command, the policy is applied on every exported route. The exported routes can also be filtered using a policy.

Policy files for this command will only recognize the following policy attributes:

- Match attributes
 - *nlri IPv6-address/mask-len*
- Action (set) attributes
 - *cost <cost>*
 - *cost-type [ase-type-1 | ase-type-2]*
 - *permit*
 - *deny*

Any other policy attribute will not be recognized and will be ignored.

The following is an example OSPFv3 export policy file:

```
entry first {
    if match any{
```

```

        nlri 2001:db8:200:300::/64;
        nlri 2001:db8:2146:23d1::/64;
        nlri 2001:db8:af31:3d0::/64;
        nlri 2001:db8:f6:2341::/64;
    } then {
        deny;
    }
}
entry second {
    if match any{
        nlri 2001:db8:304::/48;
        nlri 2001:db8:ca11::/48;
        nlri 2001:db8:da36::/48;
        nlri 2001:db8:f6a6::/48;
    } then {
        cost 220;
        cost-type ase-type-2;
        permit;
    }
}
}

```

Example

The following command enables OSPFv3 to export RIPv3-related routes and associates a policy redistrib:

```
enable ospfv3 export ripng redistrib
```

History

This command was first available in ExtremeXOS 11.2.

The tag keyword was removed in ExtremeXOS 11.4.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

enable ospfv3 restart-helper-lsa-check

```
enable ospfv3 [[vlan | tunnel] all | {vlan} vlan-name | {tunnel} tunnel-name | area area-identifier] restart-helper-lsa-check
```

Description

This command configures the restart helper router to terminate *OSPF* graceful restart when received LSAs would affect the restarting router. This will occur when the restart helper receives an LSA that will be flooded to the restarting router or when there is a changed LSA on the restarting router's retransmission list when graceful restart is initiated.

Syntax Description

vlan	<i>VLAN.</i>
all	All VLANs.
<i>vlan-name</i>	VLAN name.
area	<i>OSPFv3</i> area.
<i>area-identifier</i>	Area identifier.
restart-helper-lsa-check	Terminate graceful restart helper mode when there is a change to an LSA.

Default

Enabled.

History

This command was first available in ExtremeXOS 21.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable ospfv3 virtual-link restart-helper-lsa-check

```
enable ospfv3 virtual-link {routerid} router-identifier {area} area-identifier restart-helper-lsa-check
```

Description

This command configures the restart helper router to terminate *OSPFv3* graceful restart when received LSAs would affect the restarting router. This occurs when the restart helper receives an LSA that will be flooded to the restarting router or when there is a changed LSA on the restarting router's retransmission list when graceful restart is initiated.

Syntax Description

virtual-link	OSPFv3 virtual link.
routerid	OSPFv3 router ID.
<i>router-identifier</i>	Router ID of neighbor OSPFv3 router.
area	OSPFv3 area.
<i>area-identifier</i>	Transit area ID of virtual link.
restart-helper-lsa-check	Terminates graceful restart helper mode when there is a change to an LSA (default is enabled).

Default

Enabled.

History

This command was first available in ExtremeXOS 21.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable pim

```
enable pim {ipv4 | ipv6}
```

Syntax Description

ipv4	Specifies the IPv4 address family.
ipv6	Specifies the IPv6 address family.

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 ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

enable pim iproute sharing

```
enable pim {ipv4 | ipv6} iproute sharing
```

Description

Enables the PIM *ECMP* feature.

Syntax Description

iproute	IP Route.
sharing	Equal Cost Multipath Routing.

Default

Disabled.

Usage Guidelines

Use this feature to allow downstream PIM router to choose multiple ECMP path to source via hash from one of the following selections without affecting the existing unicast routing algorithm:

- Source
- Group
- Source-Group
- Source-Group-Next-Hop

This feature does load splitting, not load balancing, and operates on a per (S, G) and (*;G) basis, splitting the load onto the available equal cost paths by hashing according to the selection criteria defined by the user.

Make sure that IP route sharing is also enabled using `enable iproute {ipv4| ipv6} sharing`.

Example

The following command enables the PIM ECMP feature:

```
enable pim ipv4 iproute sharing
```

History

This command was first available in ExtremeXOS 15.3.2.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

enable pim snooping

```
enable pim snooping {{vlan} name}
```

Description

Enables PIM snooping globally or on one or all VLANs.

Syntax Description

<i>name</i>	Specifies a VLAN.
-------------	-------------------

Default

Disabled.

Usage Guidelines

PIM snooping does not require PIM to be enabled. However, IGMP snooping must be disabled on VLANs that use PIM snooping. PIM snooping and MVR cannot be enabled simultaneously on a switch. PIM snooping should not be enabled on a VLAN that supports PIM-DM neighbors.

Example

The following example enables PIM snooping on the default VLAN:

```
enable pim snooping default
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

NEW!

enable pim ssm tunnel

```
enable pim {ipv4} ssm tunnel [tunnel_name | tunnel all]
```

Description

Enables PIM SSM tunnels on an IP interface.

Syntax Description

ipv4	Specifies the IPv4 address family.
tunnel	Specifies the tunnel name.
<i>tunnel_name</i>	Specifies to configure PIM information for interfaces.
all	Specifies to configure PIM information on all Layer 3 interfaces.

Default

Disabled on all interfaces.

Usage Guidelines

This command enables PIM-SSM on the specified Layer 3 VLAN.

PIM-SM must also be configured on the interface for PIM to begin operating (which includes enabling IP multicast forwarding).

IGMPv3 include messages for multicast addresses in the SSM range are only processed by PIM if PIM-SSM is enabled on the interface. Any non-IGMPv3 include messages in the SSM range are not processed by PIM on any switch interface, whether SSM is enabled or not.

Example

The following example enables PIM-SSM multicast routing on tunnel accounting:

```
enable pim ssm tunnel accounting
```

History

This command was first available in ExtremeXOS 32.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

enable pim ssm vlan

```
enable pim {ipv4 | ipv6} ssm vlan [vlan_name | all]
```

Description

Enables PIM SSM on an IP interface.

Syntax Description

ipv4	Specifies the IPv4 address family.
ipv6	Specifies the IPv6 address family.
<i>vlan_name</i>	Specifies a <u>VLAN</u> name.
all	Specifies all VLANs.

Default

Disabled on all interfaces.

Usage Guidelines

This command enables PIM-SSM on the specified Layer 3 VLAN.

PIM-SM must also be configured on the interface for PIM to begin operating (which includes enabling IP multicast forwarding).

IGMPv3 include messages for multicast addresses in the SSM range are only processed by PIM if PIM-SSM is enabled on the interface. Any non-IGMPv3 include messages in the SSM range are not processed by PIM on any switch interface, whether SSM is enabled or not.

Example

The following example enables PIM-SSM multicast routing on VLAN accounting:

```
enable pim ssm vlan accounting
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

enable policy

```
enable policy
```

Description

This command enables the ONEPolicy functionality.

Syntax Description

This command has no arguments or variables.

Default

None.

Usage Guidelines

None.

Example

The following example shows how to enable ONEPolicy:

```
x450G2-48t-10G4.4 # enable policy
```

History

This command was first available in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable port

enable port [*port_list* | **all**]

Description

Enables a port.

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
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.

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 ports 3, 5, and 12 through 15 on the switch:

```
enable port 3,5,12-15
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable ports mlag-id

```
enable ports [mlag-id mlag_id]
```

Description

Enables the current ports associated with the given ID.

Syntax Description

mlag-id	Port associated with MLAG.
<i>mlag_id</i>	MLAG identifier value of the MLAG port. Range is 1–65,000.

Default

N/A.

Usage Guidelines

If any ports are added or deleted from the LAG, the port state for those ports is not changed.

In MLAG orchestration mode, this command is executed on the other MLAG peer before it is executed on the MLAG peer on which the command is run. In orchestration mode, if the MLAG port numbers are not same on both the peers, it is possible that a different set of port numbers is enabled on the different MLAG peers. This command helps ensure that the correct set of ports associated with the MLAG ID is enabled.

If the port associated with the given MLAG ID is a load shared port, all the member ports associated with this load shared group are enabled.

If the port associated with the given MLAG ID is a virtual port, the command is ignored.

Example

The following example enables the ports associated with MLAG ID "123":

```
# enable ports mlag-id 123
```

History

This command was first available in ExtremeXOS 32.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, 5720 series switches.

enable radius

```
enable radius {mgmt-access | netlogin}
```

Description

Enables the *RADIUS* client on the switch.

Syntax Description

mgmt-access	Specifies the switch management RADIUS authentication server.
netlogin	Specifies the network login RADIUS authentication server.

Default

RADIUS authentication is disabled for both switch management and network login by default.

Usage Guidelines

Before you enable RADIUS on the switch, you must configure the servers used for authentication and configure the authentication string (shared secret) used to communicate with the RADIUS authentication server.

To configure the RADIUS authentication servers, use the following command:

```
configure radius {mgmt-access | netlogin} [primary | secondary] server  
[ipaddress | hostname] {udp_port} client-ip [ipaddress] {vrvr_name}
```

To configure the shared secret, use the following command:

```
configure radius {mgmt-access | netlogin} [primary | secondary] shared-  
secret {encrypted} string
```

If you do not specify a keyword, RADIUS authentication is enabled on the switch for both management and network login. When enabled, all web, Telnet, and SSH logins are sent to the RADIUS servers for

authentication. When used with a RADIUS server that supports ExtremeXOS CLI authorization, each CLI command is sent to the RADIUS server for authorization before it is executed.

Use the `mgmt-access` keyword to enable RADIUS authentication for switch management functions.

Use the `netlogin` keyword to enable RADIUS authentication for network login.

Example

The following command enables RADIUS authentication on the switch for both management and network login:

```
enable radius
```

The following command enables RADIUS authentication on the switch for network login:

```
enable radius netlogin
```

History

This command was first available in ExtremeXOS 10.1.

The `mgmt-access` and `netlogin` keywords were added in ExtremeXOS 11.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable radius-accounting

```
enable radius-accounting {mgmt-access | netlogin}
```

Description

Enables *RADIUS* accounting.

Syntax Description

mgmt-access	Specifies the switch management RADIUS accounting server.
netlogin	Specifies the network login RADIUS accounting server.

Default

RADIUS accounting is disabled for both switch management and network login by default.

Usage Guidelines

The RADIUS client must also be enabled.

Before you enable RADIUS accounting on the switch, you must configure the servers used for accounting and configure the authentication string (shared secret) used to communicate with the RADIUS accounting server.

To configure the RADIUS accounting servers, use the following command:

```
configure radius-accounting {mgmt-access | netlogin} [primary |  
secondary] server [ipaddress | hostname] {tcp_port} client-ip [ipaddress]  
{vr vr_name}
```

To configure the shared secret, use the following command:

```
configure radius-accounting {mgmt-access | netlogin} [primary | secondary] shared-secret {encrypted}  
string
```

If you do not specify a keyword, RADIUS accounting is enabled on the switch for both management and network login.

Use the mgmt-access keyword to enable RADIUS accounting for switch management functions.

Use the netlogin keyword to enable RADIUS accounting for network login.

Example

The following command enables RADIUS accounting on the switch for both management and network login:

```
enable radius-accounting
```

The following command enables RADIUS accounting for network login:

```
enable radius-accounting netlogin
```

History

This command was first available in ExtremeXOS 10.1.

The mgmt-access and netlogin keywords were added in ExtremeXOS 11.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable radius dynamic-authorization

```
enable radius dynamic-authorization
```

Description

Enables dynamic authorization RADIUS accounting.

Syntax Description

This command has no arguments or variables.

Default

Dynamic authorization RADIUS accounting is disabled by default.

Usage Guidelines

Before you enable RADIUS on the switch, you must configure the servers used for authentication and configure the authentication string (shared secret) used to communicate with the RADIUS authentication server.

To configure the RADIUS authentication servers and shared secret, use the following command:

```
configure radius dynamic-authorization index [nas-ip [ignore | require]  
| server [host_ipaddr | host_ipV6addr | hostname] {tls {tls_port}}  
client-ip [client_ipaddr | client_ipV6addr] {vr vr_name} {shared-secret  
{encrypted} secret}
```

Example

The following command enables dynamic authorization RADIUS authentication on the switch:

```
enable radius dynamic-authorization
```

History

This command was first available in ExtremeXOS 22.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

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 is available on all platforms with a Base license or higher as described in the [***Switch Engine 32.3 Feature License Requirements***](#) document.

Platform Availability

This command is available on all platforms with an Edge, Advanced Edge, or Core license.

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

Disabled.

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 first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms with a Base license or higher as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

enable rip export

```
enable rip export [bgp | direct | e-bgp | i-bgp | ospf | ospf-
    extern1 | ospf-extern2 | ospf-inter | ospf-intra | static | isis
    | isis-level-1 | isis-level-1-external | isis-level-2 | isis-level-2-
    external ] [cost number {tag number} | policy policy-name]
```

Description

Enables *RIP* to redistribute routes from other routing functions.

Syntax Description

bgp	Specifies <i>BGP</i> routes.
direct	Specifies interface routes (only interfaces that have IP forwarding enabled are exported).
e-bgp	Specifies E-BGP routes.
i-bgp	Specifies I-BGP routes.
ospf	Specifies all <i>OSPF</i> routes.
ospf-extern1	Specifies OSPF external route type 1.
ospf-extern2	Specifies OSPF external route type 2.
ospf-inter	Specifies OSPF-inter area routes.
ospf-intra	Specifies OSPF-intra area routes.

static	Specifies static routes.
isis	Specifies IS-IS routes.
isis-level-1	Specifies ISIS Level 1 routes.
isis-level-1-external	Specifies ISIS Level 1 External routes.
isis-level-2	Specifies ISIS Level 2 routes.
isis-level-2-external	Specifies ISIS Level 2 External routes.
cost <i>number</i>	Specifies the cost metric, from 0-15. If set to 0, RIP uses the route metric obtained from the route origin.
tag <i>number</i>	Specifies a tag number.
<i>policy-name</i>	Specifies a policy.

Default

Disabled.

Usage Guidelines

This command enables the exporting of BGP, 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 RIP. If the cost metric is set to 0, the cost is inserted from the route. For example, with BGP, the cost could be the MED or the length of the BGP path. The tag value is used only by special routing applications. Use 0 if you do not have specific requirements for using a tag.

Each protocol can have a policy associated with it to control or modify the exported routes.

Example

The following command enables RIP to redistribute routes from all OSPF routes:

```
# enable rip export ospf cost 0
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms with a Base license or higher as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

enable rip originate-default cost

```
enable rip originate-default {always} cost number {tag number}
```

Description

Configures a default route to be advertised by *RIP*.

Syntax Description

always	Specifies to always advertise the default route.
<i>cost number</i>	Specifies a cost metric. The range is 1 - 15.
tag <i>number</i>	Specifies a tag number.

Default

Disabled.

Usage Guidelines

If **always** is specified, RIP always advertises the default route to its neighbors. If **always** is not specified, RIP advertises a default route only if a reachable default route is in the system route table.

The default route advertisement is filtered using the out policy.

The cost metric is inserted for all RIP-learned, static, and direct routes injected into RIP. The tag value is used only by special routing applications.

Example

The following command configures a default route to be advertised by RIP if there is a default route in the system routing table:

```
# enable rip originate-default cost 7
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms with a Base license or higher as described in the ***Switch Engine 32.3 Feature License Requirements*** document.

enable rip poisonreverse

```
enable rip poisonreverse
```

Description

Enables poison reverse algorithm for *RIP*.

Syntax Description

Enables poison reverse algorithm for RIP.

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 first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms with a Base license or higher as described in the ***Switch Engine 32.3 Feature License Requirements*** document.

enable rip splithorizon

```
enable rip splithorizon
```

Description

Enables the split horizon algorithm for *RIP*.

Syntax Description

Enables the split horizon algorithm for RIP.

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 first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms with a Base license or higher as described in the **Switch Engine 32.3 Feature License Requirements** document.

enable rip triggerupdates

Triggered updates are a mechanism for immediately notifying a router's neighbors when the router adds or deletes routes or changes their metric.

```
enable rip triggerupdates
```

Description

Enables the trigger update mechanism.

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:

```
3 enable rip triggerupdate
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms with a Base license or higher as described in the [***Switch Engine 32.3 Feature License Requirements***](#) document.

enable rip use-ip-router-alert

```
enable rip use-ip-router-alert
```

Description

Enables the router alert IP option in the outgoing [*RIP*](#) control packets.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

None.

Example

The following command enables the RIP router alert IP option:

```
# enable rip use-ip-router-alert
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms with a Base license or higher as described in the [***Switch Engine 32.3 Feature License Requirements***](#) document.

enable ripng

enable **ripng**

Description

Enables *RIPng* for the whole router.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

Although RIPng is useful in small networks, it 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.

For larger networks, consider *OSPFv3* as an alternative IGP.

Example

The following command enables RIPng for the whole router:

```
enable ripng
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Edge, Advanced Edge, or Core license. For licensing information, see the [Switch Engine 32.3 Feature License Requirements](#) document.

enable ripng export

```
enable ripng export [direct | ospfv3 | ospfv3-extern1 | ospfv3-extern2
| ospfv3-inter | ospfv3-intra | static | isis | isis-level-1 | isis-
level-1-external | isis-level-2 | isis-level-2-external | bgp | e-bgp
i-bgp] [cost number {tag number} | policy policy-name]
```

Description

Enables *RIPng* to redistribute routes from other routing functions.

Syntax Description

direct	Specifies interface routes (only interfaces that have IP forwarding enabled are exported).
ospfv3	Specifies all <i>OSPFv3</i> routes.
ospfv3-extern1	Specifies OSPFv3 external route type 1.
ospfv3-extern2	Specifies OSPFv3 external route type 2.
ospfv3-inter	Specifies OSPFv3-inter area routes.
ospfv3-intra	Specifies OSPFv3-intra area routes.
static	Specifies static routes.
isis	Specifies IS-IS routes.
isis-level-1	Specifies IS-IS Level 1 routes.
isis-level-1-external	Specifies IS-IS Level 1 External routes.
isis-level-2	Specifies IS-IS Level 2 routes.
isis-level-2-external	Specifies IS-IS Level 2 External routes.
bgp	Specifies <i>BGP</i> IPv6 routes
e-bgp	Specifies EBGp routes.
i-bgp	Specifies IBGP routes.
cost <i>number</i>	Specifies the cost metric, from 0-15. If set to 0, RIPng uses the route metric obtained from the route origin.
tag <i>number</i>	Specifies a tag number.
<i>policy-name</i>	Specifies a policy.

Default

Disabled. However, direct routes will always be advertised for all the interfaces where RIPng is enabled. For those interfaces where RIPng is not enabled, the corresponding direct route could be redistributed if direct route export is enabled through this command.

Default tag is 0.

Usage Guidelines

This command enables the exporting of static, direct, IS-IS, and OSPFv3-learned routes from the routing table into the RIPng domain. You can choose which types of IS-IS or OSPFv3 routes are injected, or you can simply choose isis or ospfv3, which will inject all learned routes (of all types) for the selected protocol.

The cost metric is inserted for all RIPng-learned, static, and direct routes injected into RIPng. If the cost metric is set to 0, the cost is inserted from the route table. The tag value is used only by special routing applications. Use 0 if you do not have specific requirements for using a tag.

Each protocol can have a policy associated with it to control or modify the exported routes. The following is sample policy file which modifies the cost of redistributed routes from OSPFv3 and statically configured routes:

```
entry filter_rt {  
    If match any {  
        Route-origin ospfv3;  
        Route-origin static;  
    }  
    then {  
        cost 10;  
    }  
}
```

Example

The following command enables RIPng to redistribute routes from all OSPFv3 routes:

```
enable ripng export ospfv3 cost 0
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Edge, Advanced Edge, or Core license. For licensing information, see the [Switch Engine 32.3 Feature License Requirements](#) document.

enable ripng originate-default

```
enable ripng originate-default {always} cost metric {tag number}
```

Description

Configures a default route to be advertised by *RIPng*.

Syntax Description

always	Specifies to advertise the default route in addition to learned default route.
cost <i>metric</i>	Specifies a cost metric. The range is 1 - 15.
tag <i>number</i>	Specifies a tag number.

Default

Disabled.

Usage Guidelines

If **always** is specified, RIPng always advertises the default route to its neighbors. If **always** is not specified, RIPng advertises a default route only if a reachable default route is in the system route table (the route is learned from other neighbors).

The default route advertisement is filtered using the out policy. Use the command, [configure ripng route-policy](#), to specify the out policy.

The cost metric is inserted for all RIPng-learned, static, and direct routes injected into RIPng. The tag value is used only by special routing applications.

Example

The following command configures a default route to be advertised by RIPng if there is a default route in the system routing table:

```
enable ripng originate-default cost 7
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Edge, Advanced Edge, or Core license. For licensing information, see the [Switch Engine 32.3 Feature License Requirements](#) document.

enable ripng poisonreverse

```
enable ripng poisonreverse {vlan vlan-name | tunnel tunnel_name | [vlan  
| tunnel] all}
```

Description

Enables the split horizon with poison reverse algorithm for *RIPng* on specified interfaces.

Syntax Description

<i>vlan-name</i>	Specifies an IPv6 configured VLAN .
<i>tunnel_name</i>	Specifies an IPv6 tunnel.
all	Specifies all interfaces.

Default

Enabled.

Usage Guidelines

Used with 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.

If both split horizon and poison reverse are enabled, poison reverse takes precedence.

Example

The following command enables split horizon with poison reverse for RIPv6 on all IPv6 interfaces in the virtual router:

```
enable ripng poisonreverse
```

The following command enables split horizon with poison reverse for all the IPv6 configured VLANs in the virtual router:

```
enable ripng poisonreverse vlan all
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Edge, Advanced Edge, or Core license. For licensing information, see the [Switch Engine 32.3 Feature License Requirements](#) document.

enable ripng splithorizon

```
enable ripng splithorizon {vlan vlan-name | tunnel tunnel_name | [vlan | tunnel] all}
```

Description

Enables the split horizon algorithm for [RIPv6](#).

Syntax Description

<i>vlan-name</i>	Specifies an IPv6 configured VLAN .
<i>tunnel_name</i>	Specifies an IPv6 tunnel.
all	Specifies all interfaces.

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 RIPv6 on all IPv6 configured interfaces:

```
enable ripng splithorizon
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Edge, Advanced Edge, or Core license. For licensing information, see the [Switch Engine 32.3 Feature License Requirements](#) document.

enable ripng triggerupdates

```
enable ripng triggerupdates {vlan vlan-name | tunnel tunnel_name | [vlan  
| tunnel] all}
```

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

<i>vlan-name</i>	Specifies an IPv6 configured VLAN .
<i>tunnel_name</i>	Specifies an IPv6 tunnel.
all	Specifies all interfaces.

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 [RIPng](#)-related traffic.

Example

The following command enables the trigger update mechanism on all IPv6 configured interfaces:

```
enable ripng triggerupdate
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Edge, Advanced Edge, or Core license. For licensing information, see the [Switch Engine 32.3 Feature License Requirements](#) document.

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 auto calibrated 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.

The switch also supports the following parameters for configuring the RMON agent, as defined in RFC2021:

- **probeCapabilities**—If you configure the probeCapabilities object, you can view the RMON MIB groups supported on at least one interface by the probe.
- **probeSoftwareRev**—If you configure the probeSoftwareRev object, you can view the current software version of the monitored device.
- **probeHardwareRev**—If you configure the probeHardwareRev object, you can view the current hardware version of the monitored device.
- **probeDateTime**—If you configure the probeDateTime object, you can view the current date and time of the probe.
- **probeResetControl**—If you configure the probeResetControl object, you can restart a managed device that is not running normally. Depending on your configuration, you can do one of the following:
 - **Warm boot**—A warm boot restarts the device using the current configuration saved in non-volatile memory.
 - **Cold boot**—A cold boot causes the device to reset the configuration parameters stored in non-volatile memory to the factory defaults and then restarts the device using the restored factory default configuration.



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 stand-alone RMON probes have traditionally been expensive. Therefore, the approach taken by Extreme Networks 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.

To view the RMON memory usage statistics for a specific memory type (for example, statistics, events, logs, history, or alarms) or for all memory types, use the following command:

```
show rmon memory {detail | memoryType}
```

Example

The following command enables the collection of RMON statistics on the switch:

```
enable rmon
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable router-discovery

```
enable router-discovery {ipv6} vlan vlan_name
```

Description

Enables router discovery advertisements on the VLAN and the processing of router discovery messages.

Syntax Description

<code>vlan_name</code>	Specifies an IPv6 configured VLAN.
------------------------	------------------------------------

Default

N/A.

Usage Guidelines

This command is only valid when the specified VLAN has an IPv6 address associated with it. After IPv6 Router Discovery is enabled on a VLAN, router advertisement messages are regularly sent on all ports associated with the VLAN.

Example

The following example enables router discovery for the VLAN "top_floor":

```
enable router-discovery vlan top_floor
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

enable sflow

```
enable sflow
```

Description

Globally enables sFlow statistical packet sampling.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command enables sFlow globally on the switch.



Note

sFlow and mirroring are not mutually exclusive. You can enable sFlow and mirroring at the same time.

Any traffic grouping using QP2 may encounter unexpected results when sFlow is enabled. For more information about [QoS](#), see the *Quality of Service* section in the [Switch Engine 32.3 User Guide](#).

Example

The following command enables sFlow sampling globally:

```
enable sflow
```


History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable sflow ports

```
enable sflow ports [ all | port_list ] { ingress | egress | both }
```

Description

Enables sFlow statistical packet sampling on a particular list of ports.

Syntax Description

<i>port_list</i>	Specifies a list of ports.
all	All ports in the system.
ingress	Enables ingress sFlow on a per-port basis.
egress	Enables egress sFlow on a per-port basis.
both	Enables both ingress and egress sFlow on a per-port basis.

Default

Ingress.

Usage Guidelines

This command enables sFlow on a particular list of ports. Ingress, egress, or a combination of both types of sampling can be enabled on a port. You also need to enable sFlow globally in order to gather statistics and send the data to the collector. Once sFlow is enabled globally, and on the ports of interest, sampling and polling begins.

Use the following command to enable sFlow globally: `enable sflow`



Note

sFlow and mirroring are not mutually exclusive. You can enable sFlow and mirroring at the same time.

For more information about mirroring, see [Configuring Slots and Ports on a Switch](#).

Example

The following command enables egress sFlow sampling on the port 3:1:

```
enable sflow ports 3:1 egress
```

History

This command was first available in ExtremeXOS 11.0.

The ingress, egress, and both keywords we added in ExtremeXOS 15.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable sharing grouping

```
enable sharing port grouping port_list {algorithm [address-based {L2 |
L3 | L3_L4 | custom} | port-based ]} {resilient-hashing [on | off]}
{distribution-mode [all | local-slot | port-lists]} {lacp | health-
check}
```

Description

Enables the switch to configure port link aggregation, or load sharing. By using link aggregation, you use multiple ports as a single logical port. Link aggregation also provides redundancy because traffic is redistributed to the remaining ports in the LAG if one port in the group goes down. LACP allows the system to dynamically configure the LAGs.

The **port-based** keyword was added to the command to support the creation of port-based load sharing groups.

Syntax Description

<i>port</i>	Specifies the master logical port for a load-sharing group or link aggregation group (LAG).
<i>port_list</i>	Specifies one or more ports or slots and ports to be grouped to the logical port.
address-based	Specifies link aggregation by address-based algorithm.
L2	Specifies address-based link aggregation by Layer 2. This is the default value.
L3	Specifies address-based link aggregation by Layer 3. Note: The L3 algorithm will be deprecated. Selection of L3 behaves the same as L3_L4. The inclusion of Layer 4 ports for distribution is not available on a per group basis. The inclusion of Layer4 ports for distribution is controlled globally for all LAGs in a switch via the <code>configure forwarding sharing [L3 L3_L4]</code> command.
L3_L4	Specifies address-based link aggregation by Layer 3 IP plus Layer4 port. Note: The inclusion of Layer4 ports for distribution is not available on a per group basis. The inclusion of Layer4 ports for distribution is controlled globally for all LAGs in a switch via the <code>configure forwarding sharing [L3 L3_L4]</code> command.

custom	Selects the custom link aggregation algorithm configured with the following command: <code>configure sharing address-based custom [ipv4 [L3-and-L4 source-only destination-only source-and-destination] hash-algorithm [xor crc-16]]</code> . The configuration of the custom option applies to all LAGs on the switch.
port-based	Supports the creation of port-based load sharing groups.
all	All active members of the group are eligible for distribution on all slots in the switch.
local-slot	If there are one or more active members of the group on the slot where traffic is received, distribution will be restricted to these local-slot members.
port-lists	If there are one or more active members of the group in the configured distribution port list for the slot on which traffic is received, distribution will be restricted to these configured ports.
resilient-hashing	Enables the resilient hashing hardware-based capability that minimizes the remapping of flows to aggregator member ports during aggregator member changes.
lACP	Specifies dynamic link aggregation, or load sharing, using the LACP.
health-check	Specifies a health check type of link aggregation group.

Default

Disabled.

Usage Guidelines

Link aggregation, or load sharing, allows you to increase bandwidth and availability between switches by using a group of ports to carry traffic in parallel between switches. The aggregation algorithm allows the switch to use multiple ports as a single logical port. For example, [VLANs](#) see the link aggregation group (LAG) as a single logical port.



Note

All ports that are designated for the LAG must be removed from all VLANs prior to configuring the LAG.

You can enable and configure dynamic link aggregation, using LACP or health-check link aggregation. Static link aggregation is the default link aggregation method.



Note

Always verify the LACP configuration by issuing the `show ports sharing` command. Look for the ports listed as being in the aggregator.

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

Link aggregation must be enabled on both ends of the link, or a network loop will result.

Any attempt to enable sharing on ports that have an configuration is denied with following error message:

ERROR: Sharing configuration on MLAG ports cannot be modified. Use "disable mlag port <port>" to remove port from MLAG group first.

**Note**

See the appropriate volume of the [Switch Engine 32.3 User Guide](#) for information on the interaction of port-based ACLs and LAGs of ports.

LAGs are defined according to the following rules:

- Although you can reference only the logical port of a LAG to a *STPD*, all the ports of a load-sharing group actually belong to the specified STPD.
- When using link aggregation, you should always reference the logical port of the LAG when configuring or viewing VLANs. VLANs configured to use other ports in the LAG will have those ports deleted from the VLAN when link aggregation becomes enabled.

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

ExtremeXOS switches use address based algorithms to determine which physical port in the LAG to use for forwarding traffic out of the switch. Refer to [configure sharing address-based custom](#) for more information on using addressing information.

For Port-based load sharing:

**Note**

If you attempt to create a port-based load sharing group with more than 16 possible aggregator ports, the following message will be displayed:

Error: The system can have a maximum of 16 ports in a load sharing group with the configured algorithm.

This message indicates enforcement of the limit of 16 aggregator ports in a port-based LAG. Existing error messages are also used to enforce the 16 aggregator port limit for port-based load sharing groups modified by the `configure sharing port add portspport_list` command.

You cannot enable sharing on ports that have MVRP enabled.

The following guidelines apply to link aggregation on all switches:

- For all switches a static LAG can contain up to 8 ports.
- An LACP LAG can include twice the number of ports as a static LAG; out of these half can be selected links and any remaining ports will be standby links.
- A Health Check LAG may contain the same number of ports as a static LAG.
- The maximum number of LAGs is 128.
- The default load-sharing algorithm is L2 address-based aggregation. Any broadcast, multicast, or unknown unicast packet is distributed across all members in the LAG.
- The available address-based parameters are L2 for Layer 2 and L3_L4 for Layer 3 plus Layer 4. If the packet is not IP, the switch applies the Layer 2 algorithm, which is the default setting.
- The custom keyword is supported on all switches.

History

This command was first available in ExtremeXOS 10.1.

The address-based algorithm was added in ExtremeXOS 11.0.

The **L2** and **L3** optional parameters were added in ExtremeXOS 11.1.

IPv6-compatibility was added in ExtremeXOS 11.2.

Dynamic link aggregation, using LACP, was added in ExtremeXOS 11.3.

The **L3_L4** optional parameter was added in ExtremeXOS 11.5.

SummitStack functionality was added in ExtremeXOS 12.0.

Health-check link aggregation was added in ExtremeXOS 12.1.3.

The **custom** keyword was added in ExtremeXOS 12.3.

The **port-based** keyword was added in ExtremeXOS 15.4.

The **resilient-hashing** keyword was added in ExtremeXOS 21.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable slpp guard

```
enable slpp guard ports [port_list | all]
```

Description

Enables the Simple Loop Protection Protocol (SLPP) Guard feature.

Syntax Description

slpp	Specifies enabling SLPP.
guard	Specifies disabling a port as soon as an SLPP PDU is received.
ports	Specifies selecting ports on which to enable SLPP guard.
<i>port_list</i>	Selects which ports to enable SLPP guard on.
all	Specifies enabling SLPP guard on all ports.

Default

By default, SLPP Guard is disabled on all ports.

Usage Guidelines

SLPP is an application that detects loops in a Split Multi-link Trunking (SMLT) network. SLPP Guard is a complementary feature that helps prevent loops in networks by administratively disabling an edge port if a switch receives an SLPP PDU from an SMLT network.

Example

The following example enables SLPP Guard on port 5:

```
# enable slpp guard ports 5
```

History

This command was first available in ExtremeXOS 30.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable smartredundancy

```
enable smartredundancy port_list
```

Description

Enables the Smart Redundancy feature on the primary port.

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
------------------	---

Default

Enabled.

Usage Guidelines

You must configure the software-controlled redundant port using the `configure ports redundant` command prior to enabling Smart Redundancy.

The Smart Redundancy feature works in concert with the software-controlled redundant port feature. With Smart Redundancy enabled on the switch, when the primary port becomes active the switch redirects all traffic to the primary port and blocks the redundant port again. If you disable Smart Redundancy, the primary port is blocked because traffic is now flowing through the redundant port.

Example

The following command enables the Smart Redundancy feature on port 4 on a switch:

```
enable smartredundancy 4
```

History

This command was available in ExtremeXOS 11.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable snmp access

```
enable snmp access {snmp-v1v2c | snmpv3}
```

Description

Selectively enables *SNMP* access on the switch.

Syntax Description

snmp-v1v2c	Specifies SNMPv1/v2c access only.
snmpv3	Specifies SNMPv3 access only.

Default

Disabled.

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 for v1/v2c/v3, provided the MIB is installed correctly on the management station. Each network manager provides its own user interface to the management facilities.

For SNMPv3, additional security keys are used to control access, so an SNMPv3 manager is required for this type of access.

This command allows you to enable either all SNMP access, no SNMP access, v1/v2c access only, or v3 access only.

To prevent any SNMP access, use the following command : `disable snmp access {snmp-v1v2c | snmpv3}`

ExtremeXOS 11.2 introduced the concept of safe defaults mode. Safe defaults mode runs an interactive script that allows you to enable or disable SNMP, Telnet, and switch ports. When you set up your switch for the first time, you must connect to the console port to access the switch. After logging in to the switch, you enter safe defaults mode. Although SNMP, Telnet, and switch ports are enabled by default, the script prompts you to confirm those settings.

If you choose to keep the default setting for SNMP—the default setting is enabled—the switch returns the following interactive script:

```
Since you have chosen less secure management methods, please remember to increase the
security of your network by taking the following actions: * change your admin password
* change your SNMP public and private strings * consider using SNMPv3 to secure network
management traffic
```

In addition, you can return to safe defaults mode by issuing the following command: `configure safe-default-script`

If you return to safe defaults mode, you must answer the questions presented during the interactive script.

For more detailed information about safe defaults mode, see the *Using Safe Defaults Mode* section in the [Switch Engine 32.3 User Guide](#).

Example

The following command enables all SNMP access for the switch:

```
enable snmp access
```

History

This command was first available in ExtremeXOS 10.1.

SNMPv3 was added to ExtremeXOS 12.2. It was also included in ExtremeXOS 11.6.4 and 12.1.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable snmp access vr

```
enable snmp access vr [vr_name | all]
```

Description

Selectively enables *SNMP* access on virtual routers.

Syntax Description

<code>vr_name</code>	Specifies the virtual router name.
all	Specifies all virtual routers.

Default

Enabled on all virtual routers.

Usage Guidelines

Use this command to enable SNMP access on any or all virtual routers.

To disable SNMP access on virtual routers, use the `disable snmp access vr` command.

To display the SNMP configuration and statistics on a specified virtual router, use the `show snmp vr_name` command.

Example

The following command enables SNMP access on the virtual router vr-finance:

```
enable snmp access vr vr-finance
```

History

This command was first available in ExtremeXOS 12.4.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable snmp community

```
enable snmp community [encrypted enc_community_name | community_name |  
                        alphanumeric-community-string | hex hex_community_name]
```

Description

Enables *SNMP* community strings.

Syntax Description

encrypted	Community name is encrypted.
<i>enc_community_name</i>	Encrypted community name.
<i>community_name</i>	Community name in ASCII format.
hex	Provide value in hexadecimal.
<i>hex_community_name</i>	Community name in hexadecimal.
<i>alphanumeric-community-string</i>	Specifies the SNMP community string name.

Default

N/A.

Usage Guidelines

This command allows the administrator to enable an snmp community that has been disabled. It sets the row status of the community to Active.

Example

The following command enables the community string named extreme:

```
enable snmp community extreme
```

History

This command was first available in ExtremeXOS 12.1.

The **hex** keyword and *hex_community_name* variable were added in ExtremeXOS 15.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable snmp notification-log

```
enable snmp notification-log [ default | name | hex hex_name | all ]
```

Description

Controls the administrative state of a log.

Syntax Description

default	The default log.
<i>name</i>	Specifies the name of the log.
hex	Provide value in hexadecimal.
<i>hex_name</i>	Name of the log in hexadecimal.
all	Specifies all logs.

Default

Disabled.

Usage Guidelines

Use this command to control the administrative state of a log.

Example

The following example enables all logs:

```
enable snmp notification-log all
```

The following example enables *nmslog2*:

```
enable snmp notification-log nmslog2
```

History

This command was first available in ExtremeXOS 15.5.

The **default** and **hex** keywords and *hex_name* variable were added in ExtremeXOS 15.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable snmp trap l3vpn

```
enable snmp trap l3vpn {vr_name}
```

Description

This command enables Layer 3 VPN MIB notification traps for the child VPN VRFs of the specified VR.

Syntax Description

<i>vr_name</i>	Specifies the name of the parent VR where this RFC 4382 scalar is applied. If <i>vr_name</i> is not provided, then this command is applied to the VR in the current context.
----------------	--

Default

Disabled.

Usage Guidelines

This command enables generation of the following Layer 3 VPN *SNMP* traps:

- mplsL3VpnVrfUp—Sent when the first IP *VLAN* becomes active and the administrative state is enabled.
- mplsL3VpnVrfDown—Sent when the last active IP VLAN becomes inactive, or the administrative state is disabled.

Example

The following example enables SNMP traps for Layer 3 VPNs on the default VR:

```
enable snmp traps l3vpn vr vr-default
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable snmp traps

```
enable snmp traps
```

Description

Turns on *SNMP* trap support.

Syntax Description

This command has no arguments or variables.

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.

To view if SNMP traps are being sent from the switch, use the [show management](#) command. The [show management](#) command displays information about the switch including the enabled/disabled state of SNMP traps being sent.

Example

The following command enables SNMP trap support on the switch:

```
enable snmp traps
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable snmp traps configuration

```
enable snmp traps configuration [save | change]
```

Description

Enables sending *SNMP* trap when saving or changing the switch configuration.

Syntax Description

configuration	Sends SNMP trap for switch configuration.
save	Generates SNMP trap when switch configuration is saved (default is disabled).
change	Generates SNMP trap when switch configuration is changed (default is disabled).

Default

The default is that SNMP traps are disabled for switch configuration changes/saves.

Example

The following example enables SNMP traps for switch configuration changes:

```
enable snmp traps configuration change
```

History

This command was first available in ExtremeXOS 22.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches..

enable snmp traps bfd

```
enable snmp traps bfd session down | session-up
```

Description

This command enables session up/down trap reception for BFD.

Syntax Description

snmp	Configure <i>SNMP</i> specific settings.
traps	Configure SNMP Trap generation settings.
bfd	BFD-specific traps.
session-down	Generate trap when BFD session goes down.
session-up	Generate trap when BFD session goes up.

Default

Both session-down and session-up.

Usage Guidelines

Use this command to enable trap reception for BFD session up/down.

Example

The following command will enable trap generation for BFD session down events.

```
# enable snmp traps bfd session-down
```

History

This command was first available in ExtremeXOS 15.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable snmp traps fdb mac-tracking

```
enable snmp traps fdb mac-tracking
```

Description

Enables *SNMP* trap generation when MAC-tracking events occur for a tracked MAC address.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

None.

Example

The following example enables SNMP traps for MAC-tracking events:

```
enable snmp traps fdb mac-tracking
```

History

This command was first available in ExtremeXOS 12.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable snmp traps identity-management

```
enable snmp traps identity-management
```

Description

Enables the identity management feature to send *SNMP* traps for low memory conditions.

Syntax Description

This command has no arguments or variables.

Default

No traps are sent.

Usage Guidelines

The low memory conditions are described in the description for the `configure identity-management stale-entry aging-time seconds` command.

Example

The following command enables the identity management SNMP trap feature:

```
enable snmp traps identity-management
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable snmp traps l2vpn

```
enable snmp traps l2vpn
```

Description

Enables *SNMP* traps associated with Layer 2 VPNs for all *MPLS* configured VLANs.

Syntax Description

This command has no arguments or variables.

Default

All Layer 2 VPN traps are disabled.

Example

The following command enables SNMP traps associated with Layer 2 VPNs:

```
enable snmp traps l2vpn
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

enable snmp traps l3vpn

```
enable snmp traps l3vpn {vr vr_name}
```

Description

Use this command to turn on *SNMP* trap support for L3 VPN.

Syntax Description

<i>vr_name</i>	Specifies the name of the parent VR where this RFC 4382 scalar is applied. If VR name is not provided, then this command is applied to the VR in the current context.
----------------	---

Default

Enabled.

Usage Guidelines

Use this command to enable generation of L3VPN SNMP traps—mplsL3VpnVrfUp and mplsL3VpnVrfDown. These trap notifications are sent under the following conditions:

- mplsL3VpnVrfUp—first IP VLAN becomes active and administrative state is enabled.
- mplsL3VpnVrfDown—last active IP VLAN becomes inactive OR administrative state is disabled.

Example

The following example enables L3 VPN SNMP traps support on the switch:

```
enable snmp traps l3vpn vr vr-default
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable snmp traps lldp

```
enable snmp traps lldp {ports [all | port_list]}
```

Description

Enables the transmission of LLDP SNMP trap notifications.

Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

Disabled.

Usage Guidelines



Note

To enable SNMP traps for LLDP MED TLVs, you must issue a separate command; use the `enable snmp traps lldp-med {ports [all | port_list]}`.

If you do not specify any ports, the system sends LLDP traps for all ports.



Note

The Avaya-Extreme proprietary TLVs do not send traps.

Example

The following command enables LLDP SNMP traps for all ports:

```
enable snmp traps lldp ports all
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable snmp traps lldp-med

```
enable snmp traps lldp-med {ports [all | port_list]}
```

Description

Enables the transmission of *LLDP SNMP* trap notifications related to LLDP MED extension TLVs.

Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

Disabled.

Usage Guidelines

If you do not specify any ports, the system sends LLDP-MED traps for all ports.

Example

The following command enables LLDP-MED SNMP traps for all ports:

```
enable snmp traps lldp-med ports all
```

History

This command was first available in ExtremeXOS 11.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable snmp traps mpls

```
enable snmp traps mpls
```

Description

Enables *SNMP* traps associated with *MPLS* for all MPLS configured VLANs.

Syntax Description

This command has no arguments or variables.

Default

All MPLS traps are disabled.

Example

The following command enables SNMP traps associated with MPLS:

```
enable snmp traps mpls
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

enable snmp traps ospf

```
enable snmp traps ospf [all | trap-map bit-map]
```

Description

Enables the *OSPF* module to send traps on various OSPF events.

Syntax Description

all	Sets RFC1850 ospfSetTrap to 0x1ffff.
trap-map	Specifies the ospfSetTrap as defined in RFC1850.
<i>bit-map</i>	Specifies the ospfSetTrap value in HEX (for example, 0x1ffff for all traps).

Default

The default is disabled.

Usage Guidelines

This command enables the OSPF module to send traps on various OSPF events.

Example

The following command sets ospfSetTrap for all traps:

```
enable snmp traps ospf all
```

History

This command was first available in ExtremeXOS 12.6.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the *Switch Engine 32.3 Feature License Requirements* document.

enable snmp traps ospfv3

```
enable snmp traps ospfv3
```

Description

Enables the transmission of *OSPFv3 SNMP* notifications.

Syntax Description

ospfv3	OSPFv3-related traps.
---------------	-----------------------

Default

The default is disabled.

Example

The following example enables the transmission of OSPFv3 SNMP notifications:

```
enable snmp traps ospfv3
```

History

This command was first available in ExtremeXOS 22.1.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

enable snmp traps port-up-down ports

```
enable snmp traps port-up-down ports [port_list | all]
```

Description

Enables port up/down trap reception for specified ports.

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
all	Specifies all ports on the switch.

Default

Enabled.

Usage Guidelines

Use this command to begin receiving [SNMP](#) trap messages when a port transitions between being up and down.

Example

The following command enables ports 3, 5, and 12 through 15 on a stand-alone switch to receive SNMP trap messages when the port goes up/down:

```
enable snmp traps port-up-down ports 3,5,12-15
```

History

This command was first available in ExtremeXOS 10.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable snmpv3

```
enable snmpv3 default-group
```

Description

Selectively enables SNMPv3 default-group access on the switch.

Syntax Description

default-group	Specifies SNMPv3 default-group.
----------------------	---------------------------------

Default

Enabled.

Usage Guidelines

This command is used to enable SNMPv3 default-group access.

Enabling SNMPv3 default-group access activates the access to an SNMPv3 default-group and the user-created SNMPv3-user part of default-group. This command produces an error if SNMPv3 access is disabled on the switch.

The default groups are: admin, initial, v1v2c_ro, v1v2c_rw.

Example

The following command enables the default group access on the switch:

```
enable snmp default-group
```

History

This command was available in ExtremeXOS 12.2.

It was also included in ExtremeXOS 11.6.4 and ExtremeXOS 12.1.2.

The default-user option was removed in ExtremeXOS 22.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable snmpv3 community

```
enable snmpv3 community [ community_index | hex hex_community_index ]
```

Description

This command enables a community entry specified by the community index.

Syntax Description

<i>community_index</i>	Community index in ASCII.
hex	Provide value in hexadecimal.
<i>hex_community_index</i>	Community index in hexadecimal.

Default

Enabled.

Usage Guidelines

This command is used to enable a community entry specified by the community index.

Example

```
enable snmpv3 community abcd
```

History

This command was available in ExtremeXOS 15.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

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 ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable ssh2

```
enable ssh2 {access-profile [access_profile | none]} {port
  tcp_port_number} {vr [vr_name | all | default]}
```

Description

Enables SSH2 server to accept incoming sessions from SSH2 clients.

Syntax Description

<i>access_profile</i>	Specifies an ACL policy.
none	Cancels a previously configured ACL policy.
port	Specifies a TCP port number. The default is port 22.
<i>vr_name</i>	Specifies a virtual router name. Note: User-created VRs are supported only on the platforms listed for this feature in the Switch Engine 32.3 Feature License Requirements document.
all	Specifies that SSH is enabled on all virtual routers.
default	Specifies that SSH is enabled on the default virtual router.

Default

The SSH2 feature is disabled by default.

Usage Guidelines

SSH2 enables the encryption of session data. You must be logged in as an administrator to enable SSH2.

Use the port option to specify a TCP port number other than the default port of 22. You can only specify ports 22 and 1024 through 65535.

Using ACLs to Control SSH Access

You can specify a list of predefined clients that are allowed SSH2 access to the switch. To do this, you configure an ACL policy to permit or deny a specific list of IP addresses and subnet masks for the SSH port. You must create an ACL policy file before you can use the access-profile option. If the ACL policy file does not exist on the switch, the switch returns an error message indicating that the file does not exist.

Use the none option to cancel a previously configured ACL.

In the ACL policy file for SSH2, the source-address field is the only supported match condition. Any other match conditions are ignored.

Policy files can also be configured using the following command:

```
configure ssh2 access-profile [ access_profile | [[addrule ] [first |  
[[before | after]previous_rule]] ] delete rule | none ]
```

Creating an ACL Policy File

To create an ACL policy file, use the `edit policy` command. For more information about creating and implementing ACL policy files, see [Policy Manager](#) and [ACLs](#).

If you attempt to implement a policy that does not exist on the switch, an error message similar to the following appears:

```
Error: Policy /config/MyAccessProfile_2.pol does not exist on file system
```

If this occurs, make sure the policy you want to implement exists on the switch. To confirm the policies on the switch, use the [ls](#) command. If the policy does not exist, create the ACL policy file.

Viewing SSH Information

To view the status of SSH2 sessions on the switch, use the `show management` command. This command displays information about the switch including the enable/disable state for SSH2 sessions and whether a valid key is present.

Example

The following command enables the SSH2 feature:

```
enable ssh2
```

The next example assumes you have already created an ACL to apply to SSH.

The following command applies the ACL `MyAccessProfile_2` to SSH:

```
enable ssh2 access-profile MyAccessProfile_2
```

History

This command was first available in the ExtremeXOS 11.0

The `access-profile` and `none` options were added in ExtremeXOS 11.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable stacking

```
enable stacking {node-address node-address}
```

Description

This command enables stacking on one or all nodes.

Syntax Description

<i>node-address</i>	Specifies the MAC address of a node in the stack. To view the MAC addresses for all nodes in a stack, enter the show stacking command.
---------------------	--

Default

Default value is stacking disabled.

Usage Guidelines

This command enables stacking on one or all nodes. When a node is operating in stacking mode, [QoS profile QP7](#) cannot be created.

For information about stacking methods, and which switches can stack with other switches, see the [Available Stacking Methods](#) topic in the [Switch Engine 32.3 User Guide](#).

If a node-address is not specified, this command first performs an analysis of the current stacking configuration on the entire stack. If the stack has not yet been configured for stacking operation, or if the configuration is self-inconsistent, the user is offered the option of invoking the easy setup function. The following message appears:

```
You have not yet configured all required stacking parameters. Would you
like to perform an easy setup for stacking operation? (y/N)
```

If you enter Yes, the easy setup procedure is invoked and you first see the following message:

```
Executing "configure stacking easy-setup" command...
```

If you enter No, the following message appears:

```
Stacking has been enabled as requested.
```

The following describes the operation performed if easy setup is neither offered nor selected.

If you do not enter any node-address, stacking is enabled on all nodes in the stack topology.

If the node-address parameter is present, stacking is enabled on the node with the specified node-address. This is the MAC address assigned to the stackable by the factory.

The [show stacking configuration](#) command shows the current configuration of this parameter as well as the value currently in use.

A node that is enabled for stacking attempts to join the active topology. If successful, it then negotiates a node role with the other nodes in the stack and becomes an operational node in the stack according to its role. The master node's configuration is applied to the node.

When this command is executed successfully, the following message appears:

```
This command will take effect at the next reboot of the specified
node(s) .
```

Example

To enable stacking on a stack:

```
# enable stacking
This command will take effect at the next reboot of the specified node(s) .
```

To enable stacking on node 5, with a MAC address 00:04:96:26:6b:ed:

```
# enable stacking node-address 00:04:96:26:6b:ed
This command will take effect at the next reboot of the specified node(s) .
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available with all licenses and platforms that support the SummitStack feature. For information about which licenses and platforms support the SummitStack feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

enable stacking-support

```
enable stacking-support
```

Description

This command enables a switch with dual-purpose hardware to participate in a stack.

Syntax Description

This command does not have additional syntax.

Default

Stacking support is disabled by default for all platforms, except for the ExtremeSwitching X450-G2 and X465 series switches, which have dedicated stacking ports that are always enabled.

Usage Guidelines

The Stacking-Support Option Control column in [Table 18](#) on page 1357 displays Yes in the rows for switch configurations for which you can enable the stacking-support option.

After you enable the stacking-support option, you must reboot the switch to activate the configuration change.

If you enable the stacking-support option on a switch and reboot, data communications on the data ports listed in [Table 18](#) on page 1357 stops, and the ports use stacking protocols instead of Ethernet protocols.

Example

To enable the stack ports, enter the following command:

```
# enable stacking-support
This setting will take effect at the next reboot of this switch.
```

History

This command was first available in ExtremeXOS 12.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable stpd

```
enable stpd {stpd_name}
```

Description

Enables the *STP* protocol for one or all STPDs.

Syntax Description

<i>stpd_name</i>	Specifies an <i>STPD</i> name on the switch.
------------------	--

Default

Enabled.

Usage Guidelines

If you want to enable the STP protocol for all STPDs, do not specify an STPD name.

Example

The following command enables an STPD named Backbone_st:

```
enable stpd backbone_st
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable stpd auto-bind

```
enable stpd stpd_name auto-bind [ {vlan} vlan_name | vlan vlan_list]
```

Description

Automatically adds ports to an *STPD* when ports are added to a member *VLAN*.

Syntax Description

<i>stpd_name</i>	Specifies an STPD name on the switch.
<i>vlan_name</i>	Specifies the name of the VLAN to have autobind enabled.
<i>vlan_list</i>	Specifies the VLAN list of IDs to have autobind enabled.

Default

The autobind feature is disabled on user-created STPDs. The autobind feature is enabled on the default VLAN that participates in the default STPD S0.

If you enable autobind and add ports to a member VLAN, those ports are automatically added to the STPD.

Usage Guidelines

If you create an STPD and a VLAN with unique names, the keywords `stpd` and `vlan` are optional.

You cannot configure the autobind feature on a network login VLAN.

In an EMISTP or PVST+ environment, when you issue this command, any port or list of ports that you add to the carrier VLAN are automatically added to the STPD with autobind enabled. In addition, any port or list of ports that you remove from a carrier VLAN are automatically removed from the STPD. This allows the STPD to increase or decrease its span as you add ports to or remove ports from a carrier VLAN.

For *MSTP*, when you issue this command, any port or list of ports that gets automatically added to an *MSTI* are automatically inherited by the CIST. In addition, any port or list of ports that you remove from an MSTI protected VLAN are automatically removed from the CIST. For more information see the section. For more information, see [Automatically Inheriting Ports--MSTP Only](#) on page 2331.

Carrier VLAN

A carrier VLAN defines the scope of the STPD, which includes the physical and logical ports that belong to the STPD and the 802.1Q tag used to transport STP BPDUs in the encapsulation mode is EMISTP or PVST+. Only one carrier VLAN can exist in a given STPD, although some of its ports can be outside the control of any STPD at the same time.



Note

The STPD ID must be identical to the VLAN ID of the carrier VLAN in that STPD.

If you configure MSTP, you do not need a carrier VLAN. With MSTP, you configure a CIST that controls the connectivity of interconnecting MSTP regions and sends BPDUs across the regions to communicate the status of MSTP regions. All VLANs participating in the MSTP region have the same privileges.

Protected VLAN

Protected VLANs are all other VLANs that are members of the STPD. These VLANs “piggyback” on the carrier VLAN. Protected VLANs do not transmit or receive *STP* BPDUs, but they are affected by STP

state changes and inherit the state of the carrier VLAN. Protected VLANs can participate in multiple STPDs, but any particular port in the VLAN can belong to only one STPD.

Enabling autobind on a protected VLAN does not expand the boundary of the STPD. However, the VLAN and port combinations are added to or removed from the STPD subject to the boundaries of the carrier VLAN.

If you configure MSTP, all member VLANs in an MSTP region are protected VLANs. These VLANs do not transmit or receive STP BPDUs, but they are affected by STP state changes communicated by the CIST to the MSTP regions. MSTIs cannot share the same protected VLAN; however, any port in a protected VLAN can belong to multiple MSTIs.

Automatically Inheriting Ports--MSTP Only

In an MSTP environment, whether you manually or automatically bind a port to an MSTI in an MSTP region, the switch automatically binds that port to the CIST. The CIST handles BPDU processing for itself and all of the MSTIs; therefore, the CIST must inherit ports from the MSTIs in order to transmit and receive BPDUs.

Displaying STP Information

To view STP configuration status of the ports on a VLAN, use the following command:

```
show {vlan} [vlan_name | vlan_list] stpd
```

Example

The examples in this section assume that you have already removed the ports from the Default VLAN.

To automatically add ports to an STPD running 802.1D, EMISTP, or PVST+ and to expand the boundary of the STPD, you must complete the following tasks:

- Create the carrier VLAN.
- Assign a VLAN ID to the carrier VLAN.
- Add ports to the carrier VLAN.
- Create an STPD (or use the default, S0).
- Enable autobind on the STPDs carrier VLAN.
- Configure the STPD tag (the STPD ID must be identical to the VLAN ID of the carrier VLAN in the STP domain).
- Enable STP.

The following example enables autobind on an STPD named s8 after creating a carrier VLAN named v5:

```
create vlan v5
configure vlan v5 tag 100
configure vlan v5 add ports 1:1-1:20 tagged
create stpd s8
enable stpd s8 auto-bind v5
configure stpd s8 tag 100
enable stpd s8
```

To automatically add ports to the CIST STPD and to expand the boundary of the STPD, you must complete the following tasks:

- Create a VLAN or use the Default VLAN. (In this example, the Default VLAN is used.)
- Create the MSTP region.
- Create the STPD to be used as the CIST, and configure the mode of operation for the STPD.
- Specify the priority for the CIST.
- Enable the CIST.

The following example enables autobind on the VLAN Default for the CIST STPD named s1. (Starting with ExtremeXOS 22.2, before configuring a user-created STP domain for MSTP, you must first disable the STPD "s0" domain, which by default is in the MSTP CIST domain, and change its operational mode to dot1d or dot1w, as only one MSTP CIST domain can be there on a switch.):

```
disable stpd s0
configure stpd s0 mode dot1d
configure mstp region 1
create stpd s1
configure stpd s1 mode mstp cist
configure stpd s1 priority 32768
enable stpd s1
```

The following example enables autobind on the VLAN math for the MSTI STPD named s2:

```
create vlan math
configure vlan math tag 2
configure vlan math add ports 2-3
configure mstp region 1
create stpd s2
configure stpd s2 mode mstp msti 1
configure stpd s2 priority 32768
enable stpd s2 auto-bind vlan math
configure stpd s2 ports link-type point-to-point 5-6
enable stpd s2
```

History

This command was first available in ExtremeXOS 10.1.

The `vlan_list` variable was added in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable stpd ports

```
enable stpd stpd_name ports [all | port_list]
```

Description

Enables the STP protocol on one or more ports.

Syntax Description

<i>stpd_name</i>	Specifies an <i>STPD</i> on the switch.
all	Specifies all ports for a given STPD.
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

Enabled.

Usage Guidelines

If you create an STPD with a unique name, the keyword `stpd` is optional.

If STP is enabled for a port, BPDUs are generated and processed on that port if STP is enabled for the associated STPD.

You must configure one or more STPDs before you can use the `enable stpd ports` command. To create an STPD, use the `create stpd stpd_name {description stpd-description}` command. If you have considerable knowledge and experience with STP, you can configure the STPD using the configure `stpd` commands. However, the default STP parameters are adequate for most networks.

Example

The following command enables slot 2, port 4 on an STPD named Backbone_st:

```
enable stpd backbone_st ports 2:4
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable stpd rapid-root-failover

```
enable stpd stpd_name rapid-root-failover
```

Description

Enables rapid root failover for faster *STP* recovery times.

Syntax Description

<i>stpd_name</i>	Specifies an <i>STPD</i> name on the switch.
------------------	--

Default

Disabled.

Usage Guidelines

This command is applicable for STPDs operating in 802.1D.

If you create an STPD with a unique name, the keyword stpd is optional.

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 ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable subvlan-proxy-arp vlan

```
enable subvlan-proxy-arp vlan [vlan-name | all]
```

Description

Enables the automatic entry of subVLAN information in the proxy ARP table.

Syntax Description

<i>vlan-name</i>	Specifies a superVLAN name.
all	Specifies all VLANs.

Default

Enabled.

Usage Guidelines

To facilitate communication between subVLANs, by default, an entry is made in the IP ARP table of the superVLAN that performs a proxy ARP function. This allows clients on one subVLAN to communicate with clients on another subVLAN. In certain circumstances, intra-subVLAN communication may not be desired for isolation reasons.



Note

The isolation option works for normal, dynamic, ARP-based client communication.

Example

The following example enables the automatic entry of subVLAN information in the proxy ARP table of the superVLAN "vsuper":

```
enable subvlan-proxy-arp vlan vsuper
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable switch bluetooth

```
enable switch bluetooth {discovery | pairing }
```

Description

Enables Bluetooth capability on a switch.

Syntax Description

switch	Designates enabling switch capabilities.
bluetooth	Designates enabling Bluetooth capabilities on a switch.
discovery	Sets discoverable mode of the switch. Default is enabled.
pairing	Sets pairing ability with other Bluetooth-capable devices. Default is enabled.

Default

By default, discovery and pairing modes are enabled.

Usage Guidelines

Using the command with no options enables Bluetooth capability on the switch. The **discovery** and **pairing** options set discoverable mode and pairing ability, respectively.

To disable Bluetooth capabilities, use the `disable switch bluetooth {discovery | pairing}` command.

To view Bluetooth and discovery/pairing status, use the `show switch bluetooth [statistics | inventory]` command.

Example

The following example enables Bluetooth capability on a switch:

```
# enable switch bluetooth
```

The following example enables discovery mode on a switch:

```
# enable switch bluetooth discovery
```

History

This command was first available in ExtremeXOS 30.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable switch locally-administered-address

```
enable switch locally-administered-address
```

Description

Directs the switch to generate locally administered per-port MAC addresses.

Syntax Description

This command has no arguments or variables.

Default

This feature is disabled by default.

Usage Guidelines

ExtremeXOS switches do not use a unique per-port MAC address when transmitting bridge protocol data units (BPDUs). As a result, switch management can become inaccessible when switch MAC addresses are learned on the wrong L2 path (corresponding to a blocking port). This command allows you to direct the switch to generate locally administered MAC addresses used by STP/RSTP/MSTP BPDUs as source MAC address instead of the switch MAC address.

Example

The following example directs the switch to generate locally administered MAC addresses:

```
enable switch locally-administered-address
```

History

This command was first available in ExtremeXOS 22.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable switch usb

```
enable switch usb
```

Description

Enables use of the switch's USB port.

Syntax Description

usb	Specifies USB port on switch.
------------	-------------------------------

Default

Enabled by default.

Usage Guidelines

This command requires a reboot to take effect. This setting persists after reboots. To remove it, use the command `disable switch usb` or use the command `unconfigure switch {all | erase [all | nvram]}` with the **all** option.

Stack support is not available. You need to run this command individually on each node in a stack.

Example

The following example enables use of the USB port:

```
enable switch usb  
This setting will take effect at the next system reboot.
```

History

This command was first available in ExtremeXOS 22.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable syslog

```
enable syslog
```

Description

Enables logging to all remote syslog host targets.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

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 `configure syslog` command.

When you use the `enable syslog` command, the exporting process of the syslog begins. This command also determines the initial state of an added remote syslog target.

Example

The following command enables logging to all remote syslog hosts:

```
enable syslog
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable tacacs

```
enable tacacs
```

Description

Enables TACACS+ authentication.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

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

Example

The following command enables TACACS+ user authentication:

```
enable tacacs
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

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 ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

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. TACACS+ authentication must also be enabled to use TACACS+ authorization. Use the following command to enable authentication:

```
enable tacacs
```

Example

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

```
enable tacacs-authorization
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable tech-support collector

```
enable tech-support collector
```

Description

Enables te tech support feature.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command turns on the tech-support feature. In the ExtremeXOS 15.4 release, the feature is disabled by default. When the feature is disabled, the previous scheduled reports are canceled, and the bootup event and critical severity events are ignored.

When the feature is enabled, if any configured collector has the report mode set to automatic, the switch automatically attempts to send switch status reports to those collectors based on the configuration setting for each individual collector.

You can always use the `run tech-support report` command to trigger a one-time report to a particular collector, or all collectors, regardless if the feature is enabled or disabled or if the collector's report mode is set to automatic or manual.

Example

The following command enables the tech-support feature:

```
enable tech-support collector
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable telnet

```
enable telnet
```

Description

Enables external 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.

ExtremeXOS 11.2 introduces the concept of safe defaults mode. Safe defaults mode runs an interactive script that allows you to enable or disable *SNMP*, Telnet, and switch ports. When you set up your switch for the first time, you must connect to the console port to access the switch. After logging in to the switch, you enter safe defaults mode. Although SNMP, Telnet, and switch ports are enabled by default, the script prompts you to confirm those settings.

If you choose to keep the default setting for Telnet—the default setting is enabled—the switch returns the following interactive script:

```
Since you have chosen less secure management methods, please remember to increase the
security of your network by taking the following actions: * change your admin password
```

* change your SNMP public and private strings * consider using SNMPv3 to secure network management traffic

In addition, you can return to safe defaults mode by issuing the following command: `configure safe-default-script`

If you return to safe defaults mode, you must answer the questions presented during the interactive script.

For more detailed information about safe defaults mode, see the *Using Safe Defaults Mode* section in the [Switch Engine 32.3 User Guide](#).

Example

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

```
enable telnet
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable tunnel

```
enable {tunnel} tunnel_name
```

Description

Allows GRE tunnels to be enabled.

Syntax Description

<i>tunnel_name</i>	GRE tunnel name
--------------------	-----------------

Default

Enabled.

Usage Guidelines

Use this command to enable GRE tunnels.

Example

```
enable myGREtunnel
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable twamp reflector

```
enable twamp reflector {restrict}
```

Description

This command enables the Session-Reflector.

Syntax Description

restrict	Restricts only TWAMP control sessions to create test sessions and reflector does not respond to TWAMP-test packets that do not match a test session created by a control session.
-----------------	---

Default

N/A.

Usage Guidelines

If the you disable the Session-Reflector, the application terminates all current TWAMP test sessions. If you specify the **restrict** keyword, only TWAMP control sessions may create test sessions and the reflector will not respond to TWAMP-test packets that do not match a test session created by a control session.

History

This command was first available in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable twamp server

```
enable twamp server
```

Description

This command enables the TWAMP server.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

History

This command was first available in ExtremeXOS 16.1.

Platform Availability

The command is available on all platforms.

enable udp-echo-server

```
enable udp-echo-server {vr vrid}{udp-port port}
```

Description

Enables UDP echo server support.

Syntax Description

<i>vrid</i>	Specifies the VR or VRF.
<i>port</i>	Specifies the UDP port.

Default

Disabled.

Usage Guidelines

UDP echo packets are used to measure the transit time for data between the transmitting and receiving ends.

Example

The following example enables UDP echo server support:

```
enable udp-echo-server
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable upm profile

```
enable upm profile profile-name
```

Description

Enables the use of the specified Universal Port profile on the switch.

Syntax Description

<i>profile-name</i>	Specifies the UPM profile to be enabled.
---------------------	--

Default

A UPM profile is enabled by default.

Example

The following command enables a UPM profile called example on the switch:

```
enable upm profile example
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the Universal Port feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

enable virtual-network remote-endpoint vxlan

```
enable virtual-network remote-endpoint vxlan [ ipaddress ipaddress {vr  
  vr_name} | all ]
```

Description

Enables a VXLAN remote endpoint.

Syntax Description

virtual-network	Virtual overlay network.
remote-endpoint	Remote tunnel endpoint information.
vxlan	VXLAN virtual networks remote endpoint.
ipaddress	Specifies an IP address of a remote endpoint.
<i>ipaddress</i>	Specifies the IP address of the desired remote endpoint.
vr	Specifies a VR/VRF instance the remote endpoint is associated with.
<i>vr_name</i>	Specifies the desired existing VR/VRF instance the remote endpoint is associated with. Default is VR-Default.
all	Specifies all remote tunnel endpoints.

Default

If a VR is not specified, VR-Default is the VR.

Usage Guidelines

Extreme Loop Recognition Protocol (ELRP) detects loops across VXLAN tunnels. If a loop is detected across the tunnel, ELRP takes down the VXLAN remote endpoint. You can use this command to re-enable the remote endpoint.

Example

The following example enables the remote endpoint at 100.1.1.1 on VR-Default (not specified, command default):

```
# enable virtual-network remote-endpoint vxlan ipaddress 100.1.1.1
```

History

This command was first available in ExtremeXOS 22.4.

Platform Availability

This command is supported on the ExtremeSwitching 5320, 5420, 5520, 5720 series switches, and stacks with 5320, 5420, 5520, 5720 slots only.

enable virtual-router

```
enable virtual-router vrf-name
```

Description

Enables a VRF.

**Note**

This command does not affect virtual routers.

Syntax Description

<i>vrf-name</i>	Specifies the name of the VR or VRF instance.
-----------------	---

Default

Enabled.

Usage Guidelines

This command is used to administratively enable or disable a VRF. The VRF specific commands are still accepted and retained by the switch. This configuration has an operational impact on the VRF.

When you enable a VRF, the software does the following:

- Enables Layer 3 protocols for the VRF.
- Marks static routes as active and adds them to the hardware forwarding tables.

Example

The following example enables VRF "vrf1":

```
enable virtual-router vrf1
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable vlan

```
enable [ {vlan} vlan_name | vlan vlan_list]
```

Description

Use this command to re-enable a VLAN that you previously disabled.

Syntax Description

<i>vlan_name</i>	Specifies the VLAN you want to enable.
<i>vlan_list</i>	Specifies the VLAN list of IDs you want to enable.

Default

Enabled.

Usage Guidelines

This command allows you to administratively enable specified VLANs that you previously disabled.

Example

The following example enables the VLAN named "accounting":

```
enable vlan accounting
```

History

This command was first available in ExtremeXOS 11.4.

The *vlan_list* option was added in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable vman cep egress filtering ports

```
enable vman cep egress filtering ports {port_list | all}
```

Description

Enables the egress filtering of frames based on their CVIDs on ports configured as CEPs.

Syntax Description

<i>port_list</i>	Specifies a list of ports.
all	Specifies all switch ports.

Default

Egress CVID filtering is disabled.

Usage Guidelines

For a given VMAN and a port configured as a CEP for that VMAN, only frames with CVIDs that have been mapped from the CEP to the VMAN are forwarded from the VMAN and out the CEP.

To view the configuration setting for the egress CVID filtering feature, use the `show ports information` command.



Note
CVID egress filtering is available only on switches that support this feature, and when this feature is enabled, it reduces the maximum number of CVIDs supported on a port. The control of CVID egress filtering applies to fast-path forwarding. When frames are forwarded through software, CVID egress filtering is always enabled.

Example

The following command enables egress CVID filtering on port 1:

```
enable vman cep egress filtering port 1
```

History

This command was first available in ExtremeXOS 12.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable vm autostart

```
enable vm vm_name autostart
```

Description

Enables automatic start-up of guest virtual machines (VMs).

Syntax Description

vm	Designates a virtual machine.
<i>vm_name</i>	Specifies the VM name.
autostart	Specifies automatic start-up of the specified VM. Default is disabled.

Default

By default, automatic start-up is disabled.

Usage Guidelines

This command enables automatically starting up a specific VM when the system starts.

You must reboot the switch for this command to take effect.

To disable automatic start-up, use the command `disable vm vm_name autostart`.

Example

The following example enables automatic start-up of VM "vm1":

```
# enable vm vm1 autostart
```

History

This command was first available in ExtremeXOS 30.3.

Platform Availability

This command is available on all platforms that support the Integrated Application Hosting (IAH) feature and have a Core or Premier license installed. For a list of platforms that support the IAH feature and for information about licenses, see the [Switch Engine 32.3 Feature License Requirements](#).

enable vm-tracking

```
enable vm-tracking
```

Description

Enables the XNV feature on the switch.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command enables the XNV feature, which tracks VMs that connect to the switch.

This command does not enable XNV on any ports. To start tracking VMs, you must enable VM tracking on one or more ports using the [enable vm-tracking ports](#) command.

Example

The following command enables the XNV feature:

```
# enable vm-tracking
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5420 and 5520 series switches.

enable vm-tracking dynamic-vlan ports

```
enable vm-tracking dynamic-vlan ports port_list
```

Description

This command enables VM-tracking dynamic VLAN on specific ports. The ALL option is not supported because VM-tracking dynamic VLAN should never be enabled on a switch's uplink port.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

Use this command to enable VM-tracking dynamic VLAN on specific ports. The ALL option is not supported because VM-tracking dynamic VLAN should not be enabled on a switch's uplink port.

Example

The following command enables VM tracking dynamic VLAN on port 2:1:

```
# enable vm-tracking dynamic-vlan ports 2:1
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on ExtremeSwitching 5420 and 5520 series switches.

enable vm-tracking ports

```
enable vm-tracking ports port_list
```

Description

Enables the XNV feature on the specified ports.

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
------------------	---

Default

Disabled.

Usage Guidelines

You must enable VM tracking on the switch with the `enable vm-tracking` command before you can use this command. This command enables VM tracking on the specified ports. You should enable VM tracking only on ports that connect directly to a server that hosts VMs that you want to track. You should never enable VM tracking on a switch uplink port.

Example

The following command enables VM tracking on port 2:1:

```
# enable vm-tracking ports 2:1
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5420 and 5520 series switches.

enable vpex

`enable` **vpex**

Description

Enables VPEx mode for using bridge port extenders (BPEs).

Syntax Description

vpex	Specifies Virtual Port Extender (VPEx).
-------------	---

Default

N/A.

Usage Guidelines

This command enables VPEX mode and allows you to refer to ports in the *slot:port* notation in applicable commands. A reboot of the switch is required for the command to take effect. After rebooting, the CLI prompt changes to show that the switch is now slot 1 (for example):

```
Slot-1 VPEX X670G2-48x-4q.14 #
```

After enabling VPEX mode, to use the BPE, you need to configure the slot assignment for the BPE, using the command: **enable vpex**

VPEX mode is not compatible with stacking mode. Only one of these modes can be enabled at a time.

Example

The following example enables VPEX mode:

```
# enable vpex
```

History

This command was first available in ExtremeXOS 22.5.

Platform Availability

This command is available on ExtremeSwitching 5420 and 5520 series switches.

enable vpex auto-configuration

```
enable vpex auto-configuration
```

Description

Enables automatic configuration of the Extended Edge Switching architecture (controlling bridge (CB) and bridge port extenders (BPEs)).

Syntax Description

vpex	Specifies Virtual Port Extender (VPEX).
auto-configuration	Specifies enabling automatic configuration of the Extended Edge Switching architecture.

Default

Disabled.

Usage Guidelines

When this command is run the controlling bridge switch detects new BPEs connected to ports not configured as cascade ports, and automatically configures cascade ports, extended slots, and LAGs/MLAGs on cascade ports.

If you want to use redundant CBs, you must create the peer relationship between redundant CBs and ensure that both CBs are up. The rest of the MLAG setup for redundant CBs is handled automatically.

To enable auto-configuration, you must first enter VPEX mode (see [enable vpex](#) on page 2353).

When auto-configuration mode is enabled, you cannot manually configure Extended Edge Switching ports using the command `configure vpex ports port_list slot slot_num`

Example

The following example enables auto-configuration mode:

```
# enable vpex auto-configuration
```

History

This command was first available in ExtremeXOS 22.6.

Platform Availability

This command is available on ExtremeSwitching 5420 and 5520 series switches.

enable vpex auto-upgrade

enable **vpex auto-upgrade**

Description

Enables automatic upgrading on Extended Edge Switching topologies.

Syntax Description

vpex	Specifies Virtual Port Extender (VPEX).
auto-upgrade	Specifies that the controlling bridge (CB) automatically upgrades bridge port extender (BPE) slots in mode (default is enabled).

Default

Automatic upgrading is enabled by default.

Usage Guidelines

Automatic upgrading can occur only when both CBs in the MLAG have the same BPE xmod versions installed, and only after all slots are synchronized between the CBs.

To enable automatic upgrading, you must first enter VPEX mode (see [enable vpex](#) on page 2353). To view the status of automatic upgrading, use the command `show vpex`.

Example

The following example enables automatic upgrading:

```
# enable vpex auto-upgrade
```

History

This command was first available in ExtremeXOS 30.5.

Platform Availability

This command is available on ExtremeSwitching 5420 and 5520 series switches.

enable vpls

```
enable vpls [vpls_name | all]
```



Note

This command has been replaced with the following command: `enable l2vpn [vpls [vpls_name | all] | vpws [vpws_name | all]]`.

This command is still supported for backward compatibility, but it will be removed from a future release, so we recommend that you start using the new command.

Description

Enables the VPLS instance specified by *vpls_name*.

Syntax Description

<i>vpls_name</i>	Identifies the VPLS within the switch (character string).
all	Specifies all VPLS.

Default

All newly created VPLS instances are enabled.

Usage Guidelines

This command enables the VPLS instance specified by *vpls_name*. By default, all newly created VPLS instances are enabled. When enabled, VPLS attempts to establish sessions between all configured peers. Services must be configured and enabled for sessions to be established successfully.

Example

The following command enables the VPLS instance myvpls:

```
enable vpls myvpls
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support *MPLS* as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

enable vpls fdb mac-withdrawal

```
enable vpls fdb mac-withdrawal
```



Note

This command has been replaced with the following command: `enable l2vpn vpls fdb mac-withdrawal`.

This command is still supported for backward compatibility, but it will be removed from a future release, so we recommend that you start using the new command.

Description

Enables the VPLS MAC address withdrawal capability.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

Use this command to enable *FDB* MAC withdrawal after it has been disabled.

Example

The following command enables MAC address withdrawal:

```
enable vpls fdb mac-withdrawal
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available only on the platforms that support *MPLS* as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

enable vpls health-check vccv

```
enable vpls vpls_name health-check vccv
```



Note

This command has been replaced with the following command: `enable l2vpn [vpls vpls_name | vpws vpws_name] health-check vccv`.

This command is still supported for backward compatibility, but it will be removed from a future release, so we recommend that you start using the new command.

Description

Enables the VCCV health check feature on the specified VPLS.

Syntax Description

<i>vpls_name</i>	Identifies the VPLS for which health check is to be enabled.
------------------	--

Default

Health check is disabled.

Usage Guidelines

Health check must be enabled on both ends of a PW to verify connectivity between two VPLS peers. Both VCCV peers negotiate capabilities at PW setup. A single VCCV session monitors a single PW. Therefore, a VPLS with multiple PWs will have multiple VCCV sessions to multiple peers.

VCCV in ExtremeXOS uses LSP ping to verify connectivity.

Example

The following command enables the health check feature on the VPLS instance myvpls:

```
enable vpls myvpls health-check vccv
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available only on the platforms that support *MPLS* as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

enable vpls service

```
enable vpls [vpls_name | all] service
```



Note This command has been replaced with the following command: `enable l2vpn [vpls [vpls_name | all] | vpws [vpws_name | all]] service`. This command is still supported for backward compatibility, but it will be removed from a future release, so we recommend that you start using the new command.

Description

Enables the configured VPLS services for the specified *vpls_name*.

Syntax Description

<i>vpls_name</i>	Identifies the VPLS within the switch (character string).
all	Specifies all VPLS.

Default

Enabled.

Usage Guidelines

This command enables the configured VPLS services for the specified *vpls_name*. When services are disabled, the VPLS is withdrawn from all peer sessions. The keyword **all** enables services for all VPLS instances.

Example

The following command enables the configured VPLS services for the specified VPLS instance:

```
enable vpls myvpls service
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support *MPLS* as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

enable vrrp group

```
enable vrrp group group_name {configuration | members}
```

Description

This command applies group configuration on individual VRs and they then become member VRs.

Syntax Description

group	Form a group of <i>VRRP</i> VRs to operate in high-scale mode.
<i>group_name</i>	Specifies the VRRP group name.
configuration	Applies group configuration on individual VRs (default).
members	Enables all VRs that are members of the group.

Default

If you do not specify, group configuration is applied to individual VRs.

Usage Guidelines

When this command is issued the primary VR of the group sends VRRP advertisements at configured intervals. Secondary VRs send at a much slower rate.

Example

The following example brings group configuration into effect on the member VRs of the group:

```
enable vrrp group ExtremeNet configuration
```

History

This command was first available in ExtremeXOS 22.2.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the VRRP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

enable vrrp vrid

```
enable vrrp {vlan [vlan_name | vlan_list] vrid [vridval | vrid_list]}
```

Description

Enables a specific [VRRP](#) instance or all VRRP instances on the switch.

Syntax Description

<i>vlan_name</i>	Specifies the name of a VRRP VLAN .
<i>vlan_list</i>	VLAN list (1-4,094).
<i>vridval</i>	Specifies the VRID for the VRRP instance to be enabled. To display the configured VRRP router instances, enter the show vrrp command.
<i>vrid_list</i>	List of virtual router IDs (1-255).

Default

N/A.

Usage Guidelines

This enables a specific VRRP instance on the device. If you do not specify a VRRP instance, all VRRP instances on this device are enabled.

Example

The following command enables all VRRP instances on the switch:

```
enable vrrp
```

History

This command was first available in ExtremeXOS 10.1.

VLAN and VR list options added in ExtremeXOS 22.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the VRRP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

enable watchdog

```
enable watchdog
```

Description

Enables the system watchdog timer.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

The watchdog timer monitors the health of the switch hardware and software events. For example, the watchdog timer reboots the switch if the system cannot reset the watchdog timer. This is caused by a long CPU processing loop, any unhandled exception, or a hardware problem with the communication channel to the watchdog. In most cases, if the watchdog timer expires, the switch captures the current CPU status and posts it to the console and the system log. In some cases, if the problem is so severe that the switch is unable to perform any action, the switch reboots without logging any system status information prior to reboot.

This command takes affect immediately.

The watchdog settings are saved in the configuration file.

To display the watchdog state of your system, use the `show switch` command.

Example

The following command enables the watchdog timer:

```
enable watchdog
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable web http

```
enable web http
```

Description

Enables hypertext transfer protocol (HTTP) access to the switch on the default HTTP port (80).

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

If HTTP access has been disabled, use this command to enable HTTP access to the switch.

Example

The following command enables HTTP on the default port:

```
enable web http
```

History

This command was first available in the ExtremeXOS 11.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable web https

```
enable web https
```

Description

Enables secure socket layer (SSL) access to the switch on the default port (443).

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

If an SSL certificate and key are not present in switch, this command will result in automatically generating both.

Use this command to allow users to connect using a more secure HTTPS connection.

To use secure HTTP access (HTTPS) for web-based login connections, you must specify HTTPS as the protocol when configuring the redirect URL. For more information about configuring the redirect URL, see the [configure netlogin redirect-page](#) command.

Example

The following command enables SSL on the default port:

```
enable web https
```

History

This command was first available in the ExtremeXOS 11.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable cli xml-mode

```
enable cli xml-mode
```

Description

Enables XML configuration mode on the switch.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command enables the XML configuration mode on the switch, however XML configuration mode is not supported for end users, and Extreme Networks strongly cautions you not to enable this mode. Use this command only under the direction of Extreme Networks.

If you inadvertently issue this command, the switch prompt will be changed by adding the text (xml) to the front of the prompt. If you see this mode indicator, please disable XML configuration mode by using the following command:

```
disable xml-mode
```

Example

The following command enables XML configuration mode on the switch:

```
enable cli xml-mode
```

History

This command was first available in an ExtremeXOS 11.2.

The **cli** keyword was added for syntax consistency in ExtremeXOS 30.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable/disable bfd vlan

```
[enable | disable] bfd vlan vlan_name
```

Description

Enables or disables BFD on a VLAN.

Syntax Description

<i>vlan_name</i>	Specifies the VLAN name.
------------------	--------------------------

Default

N/A.

Usage Guidelines

Use this command to enable or disable BFD on a VLAN.

Example

The following command enables the bfd on the VLAN named finance:

```
# enable bfd vlan finance
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

enable/disable xml-notification

```
[enable|disable] xml-notification [all | target]
```

Description

Enables or disables Web server target(s).

Syntax Description

<i>target</i>	wSpecifies the configured target.
---------------	-----------------------------------

Default

By default, the target Web server is not enabled for xml-notifications. You have to explicitly enable it, and the display value is “no.”

Usage Guidelines

Use the enable option to enable Web server target(s) in order to receive events from ExtremeXOS modules and to send out events to the targeted Web server(s).

Use the disable option to disable the Web server target(s).

Example

The following command enables all of the configured targets:

```
enable xml-notification all
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

ENDIF

ENDIF



Note

This is a script command and operates only in scripts when scripting is enabled with the following command: `enable cli scripting {permanent}`.

Description

Causes the IF construct to be terminated.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

The ENDIF command should be used after the `IF _expression THEN statement(s)` command.

You can insert comments by using a number sign (#). CLI scripting must be enabled to use this command.

Example

The following example executes the show switch command if the value of the variable is greater than 2 and execute the show vlan command otherwise:

```
IF ($x > 2) THEN
    show switch
ELSE
    show vlan
ENDIF
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

ENDWHILE

ENDWHILE



Note

This is a script command and operates only in scripts when scripting is enabled with the following command: `enable cli scripting {permanent}`.

Description

Causes the WHILE construct to be terminated.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

The ENDWHILE command must be used after a corresponding WHILE *_expression* *DO statement (s)* command.

You can insert comments by using a number sign (#). CLI scripting must be enabled to use this command.

Example

This example creates 10 VLANs, named x1 to x10:

```
set var x 1

WHILE ($x <= 10) DO

    create vlan v$x

    set var x ($x + 1)

ENDWHILE
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

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.

When you issue this command, you are asked to save your configuration changes to the current, active configuration. Enter y if you want to save your changes. Enter n if you do not want to save your changes.

Example

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

```
exit
```

A message similar to the following is displayed:

```
Do you wish to save your configuration changes to primary.cfg? (y or n)
```

Enter y if you want to save your changes. Enter n if you do not want to save your changes.

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

history

```
history {last | -c | -d start {- {end}}
```

Description

Displays a list of all the commands entered on the switch, and enables the clearing or deleting of entries.

Syntax Description

<i>last</i>	Specifies the number of most recent entries of command history to display, or all of the entries if not specified. Range is 1-2147483647.
-c	Specifies to Clear all entries of command history.
-d	Specifies to Delete an entry or a range of entries of command history.
<i>start</i>	Specifies the start of the range of command history entries (or a single entry) to delete. Range is 1-2147483647.
-	Range separator.
<i>end</i>	Specifies the end of the range of command history entries to delete. Range is 1-2147483647.

Default

N/A.

Usage Guidelines

ExtremeXOS software “remembers” all the commands you entered on the switch.

Use the history command to display a list of these commands.

Example

The following command displays all the commands entered on the switch:

```
# history
1  enable ssh
2  show switch
3  show ver images
4  show switch
5  configure timezone name EST -300 autodst name EDT
6  show switch
7  configure time 8 17 2021 14 58 0
8  save
9  show switch
10 history
```

The following command deletes the history for the range of 4 and 5 from the previous list:

```
# history -d 4 - 5
```

The following command clears the history:

```
history -c
```

If you use a command more than once, consecutively, the history will only list the first instance.

History

This command was first available in ExtremeXOS.

The ability to clear and delete entries was added in ExtremeXOS 31.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

IF ... THEN

IF (**_expression**) **THEN**



Note

This is a script command and operates only in scripts when scripting is enabled with the following command: [enable cli scripting {permanent}](#).

Description

Optionally executes a code block based on the condition supplied.

Syntax Description

expression	Specifies the condition for which the statements should be executed.
statements	Actions to be executed when the specified conditions are met.

Default

N/A.

Usage Guidelines

This command is usually followed by statements that are executed if the condition evaluates to true.

It can also be followed by an ELSE block, which is executed if the condition evaluates to false.

The IF construct should be terminated by an **ENDIF** command.

The **_expression** must be enclosed in parentheses.

The IF construct can be nested inside other IF and WHILE constructs. Nesting is supported up to five levels. If there is incorrect nesting of IF conditions, an error message is displayed. If a user tries to execute more than five nested IF conditions, an error message is displayed.

The operators mentioned in [Using Operators](#) can be used in an `_expression` in an IF condition.

You can insert comments by using a number sign (#).

Example

The following example executes the `show switch` command if the value of the variable is greater than 2 and executes the `show vlan` command otherwise:

```
IF ($x > 2) THEN
    show switch
ELSE
    show vlan
ENDIF
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

install bootrom

```
install bootrom [from-image | fname | local-file] {slot slot-number}
               {reboot}
```

Description

Installs a new version of the ExtremeXOS BootROM image.

Syntax Description

from-image	Specifies using BootROM image packaged in the booted software image.
<i>fname</i>	Specifies the BootROM image file name of previously downloaded image.
<i>local-file</i>	Specifies using a BootROM image file on a local file path (USB is <code>/usr/local/ext</code> , internal memory is <code>/usr/local/tmp</code> , and home directory is <code>/usr/local/cfg</code>).

slot	For SummitStacks, specifies installing the BootROM image on a particular node (slot).
<i>slot-number</i>	On a SummitStack, selects which node the BootROM image should be installed on.
reboot	Reboots the switch after the image is installed.

Default

N/A.

Usage Guidelines

When you download a BootROM image (`download bootrom [[ipaddress | hostname] filename {{vr} vrname} {block-size block_size} {slot slotid} {install} {reboot}}`), you are prompted to install the image immediately after the download is finished. If you choose to install the image at a later time, you can use this command to install the software on the switch.

The BootROM image file is designated with a `.xtr` file extension.

For ExtremeSwitching 5420 and 5520 series switches only, the ExtremeXOS (`.xos`) image file includes the bootROM image. The **from-exos** specifies using the BootROM version packaged with the ExtremeXOS image. You do not need to specify a file name. After the `.xos` image installation is finished, and when a new `.xos` image is in use, the bootROM image is available in ExtremeXOS file system (`/exos/bin`). When this command is run with the **from-exos** option, this image will be used for bootROM upgrade. To see what version bootROM is installed on each partition, use the `show version {detail | process name | images {partition partition}}` command.

Displaying the BootROM Versions

To display the BootROM version for the switch, use the `show version` command.

Local File Name Character Restrictions

When specifying a local or remote file name, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-)
- Underscore (_)

SummitStack Only

You can run this command only from the master node.

Example

The following example installs the previously downloaded BootROM image file `summitX-1.0.1.5-bootrom.xtr`:

```
# install bootrom summitX-1.0.1.5-bootrom.xtr
```

The following messages appear:

```
Installing bootrom...
Writing bootrom...
.....
.....
.....
Verifying Flash contents...
.....
.....
.....
.....
bootrom written.
Bootrom installed successfully
```

The following example installs the bootROM image file `uboot.xtr` located on the local file system at `/usr/local/tmp/`:

```
# install bootrom /usr/local/tmp/uboot.xtr

Downloading to Switch.
Installing bootrom...

Bootrom installed successfully.  It will be used on the next reboot.
```

History

This command was first available in ExtremeXOS 11.0.

Support for SummitStack was added in ExtremeXOS 12.0.

The *local-file* option was added in ExtremeXOS 30.7.

The **from-exos** option was removed in ExtremeXOS 31.5.

The **from-image** option was added in ExtremeXOS 31.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches and stacks.

install firmware

```
install firmware {force} {slot slot-number}
```

Description

This command upgrades the ExtremeSwitching Universal platforms using images from the installed Switch Engine package.

Syntax Description

force	Specifies that a new image is installed without a version check.
slot	Slot for firmware installation.
<i>slot-number</i>	Slot number.

Default

N/A.

Usage Guidelines

On the ExtremeSwitching 5520 switch, use the `install firmware` command to upgrade the system FPGA and PLD images.

Firmware images are bundled with ExtremeXOS software images.

On the ExtremeSwitching 5520 switch, the ExtremeXOS software automatically compares the existing firmware image flashed into the hardware with the firmware image bundled with the ExtremeXOS image. You can also use the `install firmware` command to compare the firmware images.

Before using the `install firmware` command in a stack, wait until the `show slot` command indicates the slots are operational. When the slots are operational, use the `install firmware` command.

The switch checks internal devices for a possible firmware upgrade. If the bundled firmware image is newer than the existing firmware image, the switch prompts you to confirm the upgrade.

- Enter `y` to upgrade the firmware.
- Enter `n` to cancel the firmware upgrade for the specified hardware and continue scanning for other hardware that needs to be upgraded.
- Enter `cr` to cancel the upgrade. After a firmware image upgrade, messages are sent to the log.

The new FPGA and PLD firmware overwrites the older versions flashed into the hardware. The switch always maintains a backup version in hardware in case the install is interrupted. Use the `reboot` command to reboot the switch and activate the new firmware.

During the firmware upgrade, do not cycle down or disrupt the power to the switch. If a power interruption occurs, the installed firmware may be corrupted. In this case, the switch uses a backup version, and you can run the upgrade again to install the latest version.

The switch displays status messages after you use the `install firmware` command. The output varies depending upon your platform and the software version running on your system.

During a firmware upgrade, the switch prompts you to save your configuration changes to the current, active configuration. Enter `y` to save your configuration changes to the current, active configuration. Enter `n` if you do not want to save your changes.

PoE firmware is always automatically upgraded or downgraded to match the operational code image. This configuration is not applicable to PoE firmware.

Sample Output--ExtremeSwitching 5520 Switch

The following is sample output from a ExtremeSwitching 5520 switch:

```
5520-24W # install firmware
Installing FPGA/PLD image(s) to slot 1. Do you want to continue?
(y - yes, n - no, <cr> - cancel) Yes
Installing firmware...

Firmware image has been updated successfully.
Installing firmware...

Firmware image has been updated successfully.
Installing firmware...

Firmware image has been updated successfully.

The FPGA/PLD image(s) were installed
successfully and will be activated upon the next system reboot.
```

Displaying Firmware Versions

To display the firmware version for all devices in the switch, use the `show version` command.

Example

The following example installs the newer firmware image(s):

```
# install firmware
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

install image

```
install image [inactive | filename | local-file] {partition} {slot
slotid} {reboot}
```

Description

Installs a new version of the ExtremeXOS software image.

Syntax Description

<i>filename</i>	Specifies the file name of a previously downloaded image.
<i>local-file</i>	Specifies using an image file on a local file path (USB is /usr/local/ext, internal memory is /usr/local/tmp, and home directory is /usr/local/cfg).
<i>partition</i>	Specifies primary or secondary partition, or specifies active or inactive for automatic determination: "active" "inactive" "primary" "secondary".
inactive	Copies an image from the active partition to the inactive partition. This includes the .xos image and all .xmod and .lst files.
slot	For SummitStacks, specifies installing the image on a particular node (slot).
<i>slotid</i>	On a SummitStack, selects which node the image should be installed on.
reboot	Reboots the switch after the image is installed.

Default

N/A.

Usage Guidelines

When you download a software image (`download [url url {vr vrname}] | image [active | inactive] [[hostname | ipaddress] filename [{vr} vrname] {block-size block_size} {partition} {install {reboot}}`), you are asked if you want to install the image immediately after the download is finished. If you choose to install the image at a later time, use this command to install the software on the switch.

The software image file can be an .xos file, which contains an ExtremeXOS core image, or an .xmod file, which contains additional functionality to supplement a core image.



Note

Beginning with ExtremeXOS 12.1, an ExtremeXOS core image must be installed on the alternate (non-active) partition. If you try to install on an active partition, the following error message appears: `Error: Image can only be installed to the non-active partition.`

When you install a new version of an ExtremeXOS image, the system automatically compares the currently installed bootROM image against the bootROM image contained in the new ExtremeXOS image. If the installed version is older, the system automatically upgrades to the bootROM version contained in the new ExtremeXOS image.

SummitStack Only

You can issue this command only from a master node.

Displaying the Software Image Version

To display the software image version running on the switch, use the `show version` or `show switch` commands.

Displaying the Downloaded Software Image Version.

To display a software image version that has been downloaded but not installed, use the `install image ?` command.

Local File Name Character Restrictions

When specifying a local file name, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-)
- Underscore (_)

Installing an ExtremeXOS Core Image

Install the software image on the alternate partition. You can continue to run the currently booted image, but to run the newly installed image, you need to set the boot partition with the `use image {partition} partition` command and reboot the switch.

Installing an ExtremeXOS module image

An ExtremeXOS module (`.xmod`) image has functionality that supplements a core image. You need to install a module onto an already installed core image. The version number of the core image and the module must match.

For more detailed information about a hitless upgrade, see the `download [url url {vr vrname} | image [active | inactive] [[hostname | ipaddress] filename {{vr} vrname} {block-size block_size}] {partition} {install {reboot}}` command.

Example

The following example installs the software image file `summitX440-11.5.1.2.xos` on a switch:

```
# install image summitX440-11.5.1.2.xos
```

The following example shows software images that have been downloaded, but not installed:

```
install image ?  
  
# install image ?  
<fname>      Image file name  
"summitX-12.1.0.52.xos"
```

History

This command was first available in ExtremeXOS 10.1.

The **slot** parameter was added to support SummitStack in ExtremeXOS 12.0.

The *local-file* option was added in ExtremeXOS 30.7.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

install image inactive

```
install image inactive {slot slot}
```

Description

Copies the image installed on the active partition to the inactive partition.

Syntax Description

inactive	Copy image from active partition to inactive partition. This includes the .xos image and all .xmod and .lst files.
slot	Copy image only on the specified slot. Default is to copy to all slots.
<i>slot</i>	Specifies slot number to copy image to.

Default

By default, for stacks, if a slot is not specified, the image is copied to all slots.

Usage Guidelines

Copying from active partition to inactive partition includes the .xos image and all .xmod and .lst files.

This command can act on any or all slots only from the master. If not from the master, the command can only act on its own slot.

Example

The following example copies the image on the active partition to the inactive partition:

```
# install image inactive
This will overwrite the image installed on the secondary partition with the image
installed on the primary partition.
Do you want to proceed? (y/N) Yes
Copying image to secondary partition... 100% complete.
Image installed to the secondary partition successfully.
```

History

This command was first available in ExtremeXOS 30.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

install license file

```
install license file filename {slot slot}
```

Description

Installs a license key file on ExtremeSwitching 5420 and 5520 series switches.

Syntax Description

file	Specifies providing the license key file.
<i>filename</i>	Specifies the name of the license key file. The file name must have a <code>.lic</code> extension.
slot	For stacks and Extended Edge Switching, specifies providing a slot for the license key file.
<i>slot</i>	For stacks and Extended Edge Switching, specifies the slot for the license key file.

Default

N/A.

Usage Guidelines

This command installs all of the license features contained in the selected license file to the specified slot (node).

To uninstall a license, use the command `uninstall license file filename [{revoke revocation_file} | withhold]{slot slot}`.

To uninstall a specific feature, use the command `uninstall license product product_name [revoke revocation_file | withhold] {slot slot}`

Example

The following example installs a license using the license file `mylicense.lic`:

```
# install license file mylicense.lic
```

History

This command was first available in ExtremeXOS 31.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

load script

`load script filename {arg1} {arg2} ... {arg9}`

Description

Loads (plays back) an ASCII-formatted configuration file or a user-written script file on the switch.

Syntax Description

<i>filename</i>	Specifies the user-defined name of the ASCII-formatted configuration file or a user-written script file. The script file is known as the XOS script file and uses the .xsf or .py file extension.
arg	Specifies up to nine variable values that can be specified by the user. The variables are created with the names CLI.ARGV1, CLI.ARGV2, ... CLI.ARGV9.

Default

N/A.

Usage Guidelines

Use this command to load an ASCII-formatted configuration file or a user-written script file.

Configuration File: After downloading the configuration file from the TFTP server, this command loads and restores the ASCII-formatted configuration file to the switch.

An ASCII-formatted configuration file uses the .xsf file extension, not the .cfg file extension. The .xsf file extension (known as the XOS script file) saves the XML-based configuration in an ASCII format readable by a text editor.

For more detailed information about the ASCII configuration file, including the steps involved to upload, download, and save the configuration, see the `upload configuration [hostname | ipaddress] filename {vr vr-name}` command.

User-Written Script File: After writing a script, this command executes the script and passes arguments to it. As with the configuration files, these files use the .xsf or .py file extension that is automatically added.

The command allows up to nine optional variable values to be passed to the script. These are created with the names CLI.ARGV1, CLI.ARGV2, CLI.ARGV3, ... CLI.ARGV9.

In addition, two other variables are always created. CLI.ARGV0 gives the count of the number of parameters passed, and CLI.ARGV0 contains the name of the script that is being executed.

To check the variable values use the command, `show var`.



Note

Only the .xsf and .py extensions are used. The load script command assumes a .py or .xsf extension and retries opening the file if the file cannot be found with the original specified name or no extension is provided.

Example

The following command loads the ASCII-formatted configuration named configbackup.xsf:

```
load script configbackup.xsf
```

After issuing this command, the ASCII configuration quickly scrolls across the screen. The following is an example of the type of information displayed when loading the ASCII configuration file:

```
script.meg_upload_config1.xsf.389 # enable snmp access
script.meg_upload_config1.xsf.390 # enable snmp traps
script.meg_upload_config1.xsf.391 # configure mstp region purple
script.meg_upload_config1.xsf.392 # configure mstp revision 3
script.meg_upload_config1.xsf.393 # configure mstp format 0
script.meg_upload_config1.xsf.394 # create stpd s0
```

ExtremeXOS 15.6 provided capability for Python scripting. Current Python scripting implementation allows a script to interact directly with the CLI interface for managing ExtremeXOS functionality. Python script files end in .py. The .py suffix on the *script* file name tells the `load script` command to use the Python interpreter to process the script file. Additionally, ExtremeXOS 15.6 introduced a synonym command: `run script`. This command functions exactly as `load script`.

History

This command was first available in ExtremeXOS 11.4.

Multiple arguments for user-written scripts were added in ExtremeXOS 12.1.

Scripting support for Python was added in ExtremeXOS 15.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

load var key

```
load var key key [var1 var2 ...]
```



Note

This is a script command and operates only in scripts or on the command line when scripting is enabled with the following command: `enable cli scripting {permanent}`.

Description

Imports the specified set of variables associated with a key into the current session.

Syntax Description

<i>key</i>	Specifies the key associated with the variables to be imported.
<i>var1 var2</i>	Specifies the variables to be imported. The first variable is mandatory, up to four more optional variables can be specified.

Default

N/A.

Usage Guidelines

The specified key should have created by the user. Also, the variables specified should have been saved using that key.

Attempting to use this command with a non-existent key results in an error message being displayed.

Example

The following example imports the variables “username,” “ipaddr,” and “vlan” from the key “blue:”

```
load var key blue username ipaddr vlan
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series 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.

When you issue this command, you are asked to save your configuration changes to the current, active configuration. Enter y if you want to save your changes. Enter n if you do not want to save your changes.

Example

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

```
logout
```

A message similar to the following is displayed:

```
Do you wish to save your configuration changes to primary.cfg? (y or n)
```

Enter y if you want to save your changes. Enter n if you do not want to save your changes.

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

ls

`ls file_name`

Description

Lists all configuration, policy, and if configured, core dump files in the system.

Syntax Description

<i>file_name</i>	Lists all the files that match the wildcard.
------------------	--

Default

N/A.

Usage Guidelines

When you use issue this command without any options, the output displays all of the configuration and policy files stored on the switch.

When you configure and enable the switch to send core dump (debug) information to the internal memory, specify the internal memory location `/usr/local/tmp` to display the core dump files stored internally. For more information about core dump files, see [Core Dump Files](#) on page 2385.

When you specify the *file-name* option, the output displays all of the files that fit the wildcard criteria.

Understanding the Output

Output from this command includes the following:

- The first column displays the file permission using the following ten place holders:
 - The first place holder displays - for a file.
 - The next three place holders display r for read access and w for write access permission for the file owner.
 - The following three place holders display r for read access permission for members of the file owner's group.
 - The last three place holders display r for read access for every user that is not a member of the file owner's group.
- The second column displays how many links the file has to other files or directories.
- The third column displays the file owner.
- The remaining columns display the file size, date and time the file was last modified, and the file name.

Core Dump Files

Core dump files have a `.gz` file extension. The file name format is: `core.process-name.pid.gz` where `process-name` indicates the name of the process that failed and `pid` is the numerical identifier of that process.

When the switch has not saved any debug files, no files appear. For information about configuring and sending core dump information to internal memory or a USB 2.0 storage device, see the [configure debug core-dumps](#) and [save debug tracefiles memorycard](#) commands.

For more detailed information about core dump files, see *Troubleshooting* section in the [Switch Engine 32.3 User Guide](#).

Example

The following command displays a list of all current configuration and policy files in the system:

```
ls
```

The following is sample output from this command:

```
total 424
-rw-r--r-- 1 root root 50 Jul 30 14:19 hugh.pol
-rw-r--r-- 1 root root 94256 Jul 23 14:26 hughtest.cfg
-rw-r--r-- 1 root root 100980 Sep 23 09:16 megtest.cfg
-rw-r--r-- 1 root root 35 Jun 29 06:42 newpolicy.pol
-rw-r--r-- 1 root root 100980 Sep 23 09:17 primary.cfg
-rw-r--r-- 1 root root 94256 Jun 30 17:10 roytest.cfg
```

The following command displays a list of all current configuration and policy files on a removable storage device:

```
ls /usr/local/ext
```

The following is sample output from this command:

```
-rwxr-xr-x 1 root 0 15401865 Mar 30 00:03 onie-11.2.0.13.xos
-rwxr-xr-x 1 root 0 10 Mar 31 09:41 test-1.pol
-rwxr-xr-x 1 root 0 10 Apr 4 09:15 test.pol
-rwxr-xr-x 1 root 0 10 Mar 31 09:41 test_1.pol
-rwxr-xr-x 1 root 0 223599 Mar 31 10:02 v11_1_3.cfg
```

The following command displays a list of all configuration and policy files with a file name beginning with the letter “a:”

```
(debug) BD-12804.1 # ls a*
```

Following is sample output from this command:

```
-rw-r--r-- 1 root 0 2062 Jan 6 09:11 abc
-rw-rw-rw- 1 root 0 1922 Jan 7 02:19 abc.xsf
1k-blocks Used Available Use%
16384 496 15888 3%
```

The following command displays a list of all .tgz files:

```
ls /usr/local/tmp/*.tgz
```

Following is sample output from this command:

```
-rwxr-xr-x 1 root 0 79076 Jan 6 09:47 old_traces.tgz
1k-blocks Used Available Use%
49038 110 48928 0%
```

History

This command was first available in ExtremeXOS 10.1.

The memorycard option was added in ExtremeXOS 11.0.

The internal-memory option was added in ExtremeXOS 11.4.

The file-name option was added in ExtremeXOS 12.2.

Support for USB 2.0 storage devices was added in ExtremeXOS 12.5.3.

Pathname support was added in ExtremeXOS 15.5.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

mkdir

```
mkdir directory_name
```

Description

Creates a new directory on the specified file system to relative to the current working directory.

Syntax Description

mkdir	Create a directory.
directory_name	Pathname of a directory.

Default

N/A.

Usage Guidelines

Use this command to create a new directory on the specified file system to relative to the current working directory.

History

This command was first available in ExtremeXOS 15.5

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

mrinfo

```
mrinfo {router_address} {from from_address} {timeout seconds} {multiple-  
response-timeout multi_resp_timeout} {vr vrname}
```

Description

Requests information from a multicast router.

Syntax Description

<i>router_address</i>	Specifies the unicast IP address of the router for which you want information.
<i>from_address</i>	Specifies the unicast IP address of the interface where the mrinfo request is generated.
<i>seconds</i>	Specifies a maximum time to wait for a response. The range is 1-30 seconds.
<i>multi_resp_timeout</i>	Specifies a maximum time to wait for additional responses after the first response is received. The range is 0 to 3 seconds.
<i>vrname</i>	Specifies a VR name.

Default

- router_address*: One of the local interface addresses.
- from_address*: IP address of interface from which the mrinfo query is generated.
- timeout**: 3 seconds
- multiple-response-timeout**: 1 second
- vr**: DefaultVR

Usage Guidelines

The last column of the mrinfo command output displays information in the following format:

```
[Metric/threshold/type/flags]
```

This information is described in detail in the [Syntax Description](#) on page 2388..

Table 23: mrinfo Command Display Data

Data	Description
Metric	This should always be 1 because mrinfo queries the directly connected interfaces of a device.
Threshold	This should always be 0 because the threshold feature is not supported in ExtremeXOS software.

Table 23: mrinfo Command Display Data (continued)

Data	Description
Type	The type specifies the multicast protocol type. Because the ExtremeXOS software only supports PIM, this value is always pim.
querier	The querier flag indicates that the queried router is the <i>IGMP</i> querier.
leaf	The leaf flag indicates that the IP interface has no neighbor router.
down	The down flag indicates that the interface link status is down.

Example

The following command requests information from multicast router 1.1.1.1:

```
Switch.1 # mrinfo 1.1.1.1
1.1.1.1 [Flags:PGM]
2.2.2.1 -> 2.2.2.2 [1/0/pim/querier]
1.1.1.1 -> 0.0.0.0 [1/0/pim/querier/leaf]
8.8.8.1 -> 8.8.8.4 [1/0/pim/querier]
3.3.3.1 -> 0.0.0.0 [1/0/pim/down]
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IPv4 multicast feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

mtrace

```
mtrace source src_address {destination dest_address} {group grp_address}
      {from from_address} {gateway gw_address} {timeout seconds} {maximum-
      hops number} {router-alert [include | exclude]} {vr vrname}
```

Description

Traces multicast traffic from the receiver back to the source.

Syntax Description

<i>src_address</i>	Specifies the unicast IP address of the multicast source.
<i>dest_address</i>	Specifies the unicast IP address of the multicast group receiver.
<i>grp_address</i>	Specifies the multicast IP address of the group.

<i>from_address</i>	Specifies the unicast IP address of the interface where the mtrace request originates. This is used as the IP destination address of the mtrace response packet.
<i>gw_address</i>	Specifies the gateway router IP address of the multicast router to which the unicast mtrace query is sent.
<i>seconds</i>	Specifies a maximum time to wait for the mtrace response before making the next attempt. The range is 1–30 seconds.
<i>number</i>	Specifies the maximum number of hops for the trace. The range is 1 to 255.
router-alert	Specifies whether the router-alert option is included or excluded in mtrace packets.
<i>vrname</i>	Specifies a VR name.

Default

destination: IP address of interface from which mtrace query is generated.

group: 0.0.0.0

from: IP address of interface from which mtrace query is generated.

gateway: 224.0.0.2 when the destination is in the same subnet as one of the IP interfaces. For a non-local destination address, it is mandatory to provide a valid multicast router address.

timeout: 3 seconds

maximum-hops: 32

router-alert: include

vr: DefaultVR

Usage Guidelines

The multicast traceroute initiator node generates a multicast query and waits for timeout period to expire. If there is no response for the timeout period, the initiator node makes two more attempts. If no response is received after three attempts, the initiator node moves to a hop-by-hop trace by manipulating the maximum hop fields to perform a linear search.

The multicast trace response data contains the following fields:

- Incoming interface address—Interface on which traffic is expected from the specific source and group
- Outgoing interface address—Interface on which traffic is forwarded from the specified source and group towards the destination
- Previous hop router address
- Input packet count on incoming interface
- Output packet count on outgoing interface
- Total number of packets for this source-group pair

- Multicast routing protocol
- Forwarding code

Extreme Networks switches set the packet count statistics field to 0xffffffff to indicate that this field is not supported.

The last column of the mtrace command output displays forwarding codes, which are described in the following table.

Table 24: mtrace Command Forwarding Codes

Forwarding Code	Description
Wrong interface	mtrace request arrived on an interface to which this router would not forward for this source and group.
Prune sent upstream	This router has sent a prune request upstream for the source and group in the mtrace request.
Output pruned	This router has stopped forwarding for this source and group in response to a prune request from the next hop router.
Hit scope boundary	The group is subject to administrative scoping at this hop.
No route	This router has no route for the source or group and no way to determine a potential route.
Wrong Last Hop	This router is not the proper last-hop router.
Not forwarding ²	This router is not forwarding for this source and group on the outgoing interface for an unspecified reason.
Reached RP/Core	Reached rendezvous point or core.
RPF Interface	mtrace request arrived on the expected RPF interface (upstream interface) for this source and group.
Multicast disabled	mtrace request arrived on an interface which is not enabled for multicast.
Info. Hidden ²	One or more hops have been hidden from this trace.
No space in packet	There was not enough room to insert another response data block in the packet.
Next router no mtrace ²	The previous hop router does not understand mtrace requests.
Admin. Prohibited ^a	mtrace is administratively prohibited.

Example

The following command initiates an mtrace for group 225.1.1.1 at IP address 1.1.1.100:

```
Switch.6 # mtrace source 1.1.1.100 group 225.1.1.1
Mtrace from 1.1.1.100 to Self via 225.1.1.1
0          34.2.2.4
-1         34.2.2.4  PIM thresh^ 0          1.1.1.100/32  RPF Interface
-2         34.2.2.3  PIM thresh^ 0          1.1.1.100/32
-3         23.1.1.2  PIM thresh^ 0          1.1.1.100/32
-4         2.2.2.1   PIM thresh^ 0          1.1.1.100/32
Round trip time 9 ms; total ttl of 4 required.
```

² ExtremeXOS switches along the mtrace path do not provide this forwarding code.

History

This command was first available in ExtremeXOS 12.4.

The **router-alert** option was added in ExtremeXOS 12.5.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IPv4 multicast feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

mv

```
mv old_name new_name
```

Description

Moves a file from the specified file system or relative to the current working directory to another file on the specified file system or relative to the current working directory.

Syntax Description

<i>old_name</i>	Specifies the current name of the configuration or policy file on the system.
<i>new_name</i>	Specifies the new name of the configuration or policy file on the system.

Default

N/A.

Usage Guidelines

Use this command to move a file from the specified file system or relative to the current working directory to another file on the specified file system or relative to the current working directory. This command provides the functionality to relocate an existing file by creating a new entry in the file system, linking the content of the existing file to the new one and removing the old entry. If given a different name, the new file can be created in the same directory as the existing file

- XML-formatted configuration files have the `.cfg` file extension. The switch only runs `.cfg` files.
- ASCII-formatted configuration files have the `.xsf` file extensions. For more information, see the [Software Upgrade and Boot Options](#) section in the [Switch Engine 32.3 User Guide](#).
- Policy files have the `.pol` file extension.
- Core dump files have the `.gz` file extension. For more information, see the [Internal Memory and Core Dump Files](#) section in the [Switch Engine 32.3 User Guide](#).

Make sure the renamed file uses the same file extension as the original file. If you change the file extensions, the file may be unrecognized by the system. For example, if you have an existing configuration file named `test.cfg`, the new file name must include the `.cfg` file extension.

You cannot rename an active configuration file (the configuration currently selected to boot the switch). To verify the configuration that you are currently using, run the `show switch {detail}` command. If you attempt to rename the active configuration file, the switch displays a message similar to the following:

```
Error: Cannot rename current selected active configuration file.
```

When you rename a file, the switch displays a message similar to the following:

```
Rename config test.cfg to config megtest.cfg on switch? (y/n)
```

Type `y` to rename the file on your system. Type `n` to cancel this process and keep the existing file name.

Case-Sensitive File Names

File names are case-sensitive. For example, if you have a configuration file named `Test.cfg`, and you attempt to rename the file with the incorrect case, for example `test.cfg`, the switch displays a message similar to the following:

```
Error: mv: unable to rename `/config/test.cfg': No such file or directory
```

Since the switch is unable to locate `test.cfg`, the file is not renamed.

Local File Name Character Restrictions

When specifying a local file name, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z).
- Numerals (0-9).
- Period (.).
- Dash (-).
- Underscore (_).

Internal Memory and Core Dump Files

Core dump files have a `.gz` file extension. The file name format is: `core.process-name.pid.gz`, where `process-name` indicates the name of the process that failed and `pid` is the numerical identifier of that process.

When you configure the switch to send core dump (debug) information to internal memory, specify the file path `/usr/local/tmp` to rename an existing core dump file. If you have a switch with a USB storage device installed, you can move and rename the core dump file to that location.

For information about configuring and saving core dump information, see the `configure debug core-dumps [off | directory_path]` and `save debug tracefiles directory_path` commands.

Example

The following command renames the configuration file named `Testb91.cfg` to `Activeb91.cfg`:

```
# mv Testb91.cfg Activeb91.cfg
```

The following example moves the configuration file named `test1.cfg` from the switch to the USB storage device:

```
# mv test1.cfg /usr/local/ext/test1.cfg
```

The following example moves the policy file named `bgp.pol` from the USB storage device to the switch:

```
# mv /usr/local/ext/bgp.pol bgp.pol
```

History

This command was first available in ExtremeXOS 10.1.

The **memorycard** option was added in ExtremeXOS 11.1.

The **internal-memory** option was added in ExtremeXOS 11.4.

Support for USB 2.0 storage devices was added in ExtremeXOS 12.5.3.

Path name support was added in ExtremeXOS 15.5.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

nslookup

```
nslookup {IPv4 | IPv6} hostname
```

Description

Displays the IP address of the requested host.

Syntax Description

IPv4	Lookup only IPv4 address(es).
IPv6	Lookup only IPv6 address(es).
<i>hostname</i>	Specifies the hostname.

Default

Lookup both IPv4 and IPv6 addresses.

Usage Guidelines

For nslookup to work, you must configure the DNS client, and the switch must be able to reach the DNS server.

By default, the command looks for both IPv4 and IPv6 addresses and reports an error only when neither an IPv4 address nor an IPv6 address is found for the host.

If the IPv4 or IPv6 option is specified, DNS lookup happens only for that address type, and an error is reported when no address of that type is found.

Host Name and Remote IP Address Character Restrictions

This section provides information about the characters supported by the switch for host names and remote IP addresses.

When specifying a host name or remote IP address, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z).
- Numerals (0-9).
- Period (.).
- Dash (-) Permitted only for host names.
- Underscore (_) Permitted only for host names.
- Colon (:).

When naming or configuring an IP address for your network server, remember the requirements listed above.

Example

The following command looks up the IP addresses of a computer with the name myhost.mydomain that has 2 IPv4 addresses and 1 IPv6 address:

```
nslookup myhost.mydomain
```

The following is sample output from the command on a switch:

```
Host "myhost.mydomain" has the IPv4 address 192.168.1.1
Host "myhost.mydomain" has the IPv4 address 192.168.1.2
Host "myhost.mydomain" has the IPv6 address 2000::1
```

History

This command was first available in ExtremeXOS 10.1.

Support for using an IP address to obtain the name of the host was added in ExtremeXOS 11.0. Support for looking up IPv6 addresses was added in ExtremeXOS 12.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

open vm console

```
open vm vm_name {console}
```

Description

Opens a session to the serial console of a virtual machine (VM).

Syntax Description

vm	Designates a virtual machine.
<i>vm_name</i>	Specifies the VM name to use with a serial console.
console	Open VM serial console (default).

Default

By default, a serial console is opened.

The VM must be internally configured to enable the serial console.

Usage Guidelines

You can disconnect the console session by typing CTRL + Y, or, if using Telnet recursively with an appropriate client, by typing CTRL +] followed by "send escape". A maximum of one session can be active for a VM.

You cannot access the serial console before starting a VM. You must start the VM, and then reboot it to gain serial console access.

The Integrated Application Hosting (IAH) feature requires the Solid State Storage Device SSD-120.

Example

The following example opens a serial console session with VM "vm1":

```
# open vm vm1 console
```

History

This command was first available in ExtremeXOS 30.3.

Platform Availability

This command is available on all platforms that support the Integrated Application Hosting (IAH) feature and have a Core or Premier license installed. For a list of platforms that support the IAH feature and for information about licenses, see the [Switch Engine 32.3 Feature License Requirements](#).

ping

```
ping {count count {start-size start-size} | continuous {start-size
start-size} | {start-size start-size {end-size end-size}}} {udp}
{dont-fragment} {ttl ttl} {tos tos} {interval interval} {vr vrid}
{ipv4 host | ipv6 host} {from} {with record-route}
```

Description

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

Syntax Description

<i>count</i>	Specifies the number of ping requests to send.
<i>start-size</i>	Specifies the size, in bytes, of the packet to be sent, or the starting size if incremental packets are to be sent.
continuous	Specifies that UDP or ICMP echo messages to be sent continuously. This option can be interrupted by pressing [Ctrl] + C .
<i>end-size</i>	Specifies an end size for packets to be sent.
udp	Specifies that the ping request should use UDP instead of ICMP.
dont-fragment	Sets the IP to not fragment the bit.
ttl	Sets the TTL value.
<i>tos</i>	Sets the TOS value.
<i>interval</i>	Sets the time interval between sending out ping requests.
vr	Specifies the virtual route to use for sending out the echo message. If not specified, <i>VR-Default</i> is used. Note: User-created VRs are supported only on the platforms listed for this feature in the Switch Engine 32.3 Feature License Requirements document.
ipv4	Specifies IPv4 transport.
ipv6	Specifies IPv6 transport. Note: If you are contacting an IPv6 link local address, you must specify the <i>VLAN</i> you are sending the message from: ping ipv6 <i>link-local address %vlan_name host</i> .
<i>host</i>	Specifies a host name or IP address (either v4 or v6).

from	Uses the specified source address. If not specified, the address of the transmitting interface is used.
with record-route	Sets the traceroute information.

Default

N/A.

Usage Guidelines

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

You use the `ipv6` variable to ping an IPv6 host by generating an ICMPv6 echo request message and sending the message to the specified address. If you are contacting an IPv6 link local address, you must specify the VLAN you are sending the message from, as shown in the following example (you must include the % sign):

```
ping ipv6 link-local address %vlan_name host.
```

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

When the IPv6 host ping fails, the following error message appears:

```
Error: cannot determine outgoing interface. Link local address must be of form LLA%  
vlan_name.
```

Due to upgrading ExtremeXOS 30.1 to 4.14 Linux kernel, ping success to local IP addresses does not depend on link-layer status. Earlier releases of ExtremeXOS had customized Linux behavior that meant that pinging a local VLAN interface would fail when the local interface was down. However, in ExtremeXOS 30.1, pinging a local VLAN interface that is down will result in a successful ping.

If you have an asymmetric routing to a specific destination (where the traffic from the source to the destination uses one path, and the return traffic uses another), use this command with option 7, **with record-route**.

For example:

```
ping <destination address> with record-route  
ping <destination address> from <source-address> with record-route
```

For more information about this option, see [enable ip-option record-route](#).

Example

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

```
ping continuous 123.45.67.8
```

The following example uses the **with record-route** option:

```
ping 10.2.1.1 from 10.2.1.2 with record-route  
Ping(ICMP) 10.2.1.1: 4 packets, 8 data bytes, interval 1 second(s).
```

```
16 bytes from 10.2.1.1: icmp_seq=0 ttl=63 time=10 ms
RR: 10.2.1.2    10.2.0.21    10.2.1.1    10.2.1.1    10.2.0.1    10.2.1.2
```

History

This command was first available in ExtremeXOS 10.1.

The **ipv6** option was added in ExtremeXOS 11.2.

The IPv6 error message was modified in ExtremeXOS 15.2.

Ping success to local IP addresses not depend on link-layer status added in ExtremeXOS 30.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

ping mac port

The ping, or loopback message (LBM), goes from the MEP configured on the port toward the given MAC address.

```
ping mac mac port port {domain} domain_name {association}
      association_name
```

Description

Allows you to ping on the Layer 2 level throughout the specified domain and MA.

Syntax Description

<i>mac</i>	Enter the unique system MAC address on the device you want to reach. Enter this value in the format XX:XX:XX:XX:XX:XX.
<i>port</i>	Enter the port number of the MEP from which you are issuing the ping.
domain	Enter this keyword.
<i>domain_name</i>	Enter the name of the domain from which you are issuing the ping.
association	Enter this keyword.
<i>association_name</i>	Enter the name of the association from which you are issuing the ping.

Default

N/A.

Usage Guidelines

You must have CFM parameters configured prior to issuing a Layer 2 ping.

In order to send a Layer 2 ping, you must specify the port (MEP), the domain, and the MA from which you are issuing the ping. An UP MEP sends the ping to all ports (except the sending port) on the VLAN that is assigned to the specified MA, and a DOWN MEP sends the ping out from that port from that MA toward the specified MAC address.

All MIPs along the way forward the LBM to the destination. The destination MP responds back to the originator with a loopback reply (LBR).

This command sends out a ping from the MEP configured on the specified port toward the specified MAC address. If you attempt to send a ping message from a port that is not configured as a MEP, the system returns an error message. If the specified MAC address is not present in the Layer 2 forwarding table (FDB), the system cannot send the ping (applies to UpMEP, not DownMEP).

Example

The following command sends a Layer 2 ping to the unique system MAC address 00:04:96:1F:A4:31 from the previously configured UP MEP (port 2:4) in the speed association in the atlanta domain:

```
ping mac 00:04:96:1F:A4:31 port 2:4 atlanta speed
```

The following is sample output from the Layer 2 ping command:

```
BD-12802.48 # ping mac 00:04:96:1e:14:70 port 2:12 "extreme" 100
Send L2 Ping from Down MEP on 2:12, waiting for responses [press Ctrl-C to abort].
42 bytes from 00:04:96:1e:14:70, seq=4 time=17 ms
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

ping mpls lsp

```
ping mpls lsp [lsp_name | any host | prefix ipNetmask] {reply-mode [ip
| ip-router-alert]} {continuous | count count} {interval interval}
{start-size start-size {end-size end-size}} {ttl ttl} {{from from}
{next-hop hopaddress}}
```

Description

Sends an MPLS ping packet to a FEC over an LSP.

Syntax Description

<i>lsp_name</i>	Specifies the LSP on which to send the MPLS echo request.
any	Allows the echo request to be sent over any available LSP.
<i>host</i>	Specifies the FEC using an ipaddress or hostname.
prefix	Specifies a prefix.
<i>ipNetmask</i>	Specifies the prefix address.
reply-mode	Specifies the return path for the MPLS echo response.
ip	Requests an IP UDP reply packet. This is the default mode.
ip-router-alert	Requests an IP UDP reply packet with the IP Router Alert option. If the reply is sent in an LSP, the router-alert label is inserted at the top of the label stack.
continuous	Sends pings continuously until the user intervenes.
<i>count</i>	Determines whether the size of the packet increments by one byte for each new MPLS echo request sent.
<i>interval</i>	Specifies the time interval (in seconds) between pings.
<i>start-size</i>	The number of payload data bytes in the MPLS ping packet. The range is from 1 - 1518 (if jumbo frames are disabled) and from 1 - the configured jumbo packet size (if jumbo frames are enabled). The default is 8 bytes.
<i>end-size</i>	Specifies that the size of the packet increments by one byte for each new MPLS echo request sent, up to the specified maximum size for the MPLS ping packet.
<i>ttl</i>	Sets the time-to-live value in the ping packet
<i>from</i>	Specifies the source IP address of the packet.
<i>hopaddress</i>	Specifies the next-hop address.

Default

Destination IP address for MPLS echo request - random, from the 127 and 128 IP address space
 IP TTL - 1
 TTL value in MPLS echo request - 255
 Destination UDP port - 3503
 Payload data packet size - 8 bytes
 Number of pings sent - 4

Usage Guidelines

This command sends an MPLS ping packet to a FEC over an LSP. The ping command, with **mpls** keyword option, can be used to verify data plane connectivity across an LSP. This is useful because not all failures can be detected using the MPLS control plane. The **lsp** keyword and *lsp_name* parameter may be used to specify the LSP on which to send the MPLS echo request. The **lsp** keyword along with the **any** keyword allows the echo request to be sent over any available LSP that terminates at host, specified as an ipaddress or hostname. If no LSP exists to the specified host, the ping command fails even though an IP routed path may exist. If the optional **next-hop** is specified, the MPLS echo request is sent along the LSP that traverses the specified node. This option is useful for specifying an LSP when multiple LSPs exist to the specified FEC. For RSVP-TE LSPs, the FEC is implied from the LSP configuration. The TTL value in the MPLS Echo Request is set to 255.

By default, the destination IP address of the MPLS echo request is randomly chosen from the 127/8 IP address space and the IP TTL is set to 1. The destination UDP port is 3503 and the sender chooses the source UDP port.

The optional **start-size** keyword specifies the number of bytes to include as payload data in the MPLS ping packet. If no *start-size* parameter is specified, the size of the payload data is eight bytes. The minimum valid *start-size* value is one. The maximum *start-size* value is variable, depending on the type of MPLS ping packet sent, but the total size of the MPLS ping packet cannot exceed the configured jumbo packet size, if jumbo frames are enabled, or 1518 if jumbo frames are disabled. If the **end-size** keyword is specified, the size of the packet increments by one byte for each new MPLS echo request sent. The next MPLS echo request is not sent until the MPLS echo response for the previous packet is received. This is useful for detecting interface MTU mismatch configurations between LSRs. The switch ceases sending MPLS echo requests when the specified *end-size* value is reached, the MPLS ping is user interrupted, or an MPLS echo response is not received after four successive retries.

The optional **reply-mode** keyword is used to specify the reply mode for the MPLS echo response. When the **ip** option is specified, the MPLS echo reply is routed back to the sender in a normal IPv4 packet. When the **ip-router-alert** option is specified, the MPLS echo reply is routed back to the sender in an IPv4 packet with the Router Alert IP option set. Additionally, if the **ip-router-alert** option is specified and the reply route is through an LSP, the Router Alert Label is pushed onto the top of the label stack. If the reply-mode is not specified, the **reply-mode ip** option applies.

Example

The following example shows a ping command and the resulting display:

```
ping mpls lsp prefix 11.100.100.212/32
Ping(MPLS) : 4 packets, 8 data bytes, interval 1 second(s).
98 bytes from 11.100.100.212: mpls_seq=0 ttl=64 time=6.688 ms
98 bytes from 11.100.100.212: mpls_seq=1 ttl=64 time=6.036 ms
98 bytes from 11.100.100.212: mpls_seq=2 ttl=64 time=6.218 ms
98 bytes from 11.100.100.212: mpls_seq=3 ttl=64 time=6.467 ms
--- ping statistics ---
4 packets transmitted, 4 received, 0% loss
round-trip min/avg/max = 6/6/6/ms
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

pwd

pwd

Description

Prints the full pathname of the current working directory.

Syntax Description

pwd	Print current working directory.
-----	----------------------------------

Default

N/A.

Usage Guidelines

Use this command to print the full pathname of the current working directory.

History

This command was first available in ExtremeXOS 15.5

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

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.

When you issue this command, you are asked to save your configuration changes to the current, active configuration. Enter y if you want to save your changes. Enter n if you do not want to save your changes.

Example

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

```
quit
```

A message similar to the following is displayed:

```
Do you wish to save your configuration changes to primary.cfg? (y or n)
```

Enter y if you want to save your changes. Enter n if you do not want to save your changes.

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

reboot

```
reboot {[time mon day year hour min sec] | cancel} {slot slot-number}  
      | node-address node-address | stack-topology {as-standby} | all} |  
      rolling}
```

Description

Reboots the switch, bridge port extenders (BPEs), or SummitStack in the specified slot at a specified date and time.

Syntax Description

time	Specifies a reboot date in mm dd yyyy format and reboot time in hh mm ss format.
cancel	Cancels a previously scheduled reboot.
slot <i>slot-number</i>	Specifies the slot number currently being used by the active stack node or BPE that is to be rebooted.
all	Specifies rebooting all attached BPEs and the controlling bridge switch. Using this option requires a Core License or above.
rolling	Specifies to reboot or upgrade a stack by rebooting one node at a time.

Default

N/A.

Usage Guidelines

If you do not specify a reboot time, the switch reboots immediately following the command, and any previously scheduled reboots are cancelled.

Prior to rebooting, the switch returns the following message:

```
Do you want to save configuration changes to primary and reboot? (y - save and reboot, n -  
reboot without save, <cr> - cancel command)
```

To cancel a previously scheduled reboot, use the cancel option.

After selecting the **rolling** option, if a new image was not installed, a rolling reboot of all the nodes will be performed.

SummitStack Only

The `reboot` command used without any parameters on the master node reboots all members of the same active topology to which the master node belongs.

This version can only be used on the master node.

The `reboot slot slot-number` command can be used on any active node. The command will reboot the active node that is currently using the specified slot number in the same active topology as the issuing node. This variation cannot be used on a node that is not in stacking mode.

The `reboot node-address node-address` command can be used on any node whether or not the node is in stacking mode. It will reboot the node whose MAC address is supplied.

The `reboot stack-topology {as-standby}` command reboots every node in the stack topology. The command can be issued from any node whether or not the node is in stacking mode. If the `as-standby` option is used, every node in the stack topology restarts with master-capability disabled. This option is useful when manually resolving a dual master situation.

The `reboot rolling` command initiates a Stack Rolling Software Upgrade (Rolling Upgrade), which reboots each node in a stack one-by-one, allowing the stack to continue functioning during the reboot. The rolling upgrade process reboots all the Standby nodes, reboots the Backup node, and then initiates a failover to make the old Backup node the new Primary. The failover results in the old Primary node rebooting. Since the selected software version on each node is the newer, upgraded software version, after each node reboots, they will run the new software version. Performing a Rolling Upgrade does not change the current method of installing software, only how the stack is rebooted.

Bridge Port Extenders (BPEs)

Under normal circumstances, it is not necessary to reboot the BPEs slots. After rebooting a controlling switch, BPE upstream port down events cause the BPE's software to bring down all extended ports. This makes the BPE slot appear as down to any adjacently attached devices, and traffic properly re-converges on any redundant paths. When the controlling bridge switch comes back up, the BPE comes back up without any intervention. Therefore, using the commands `reboot` or `unconfigure switch {all | erase [all | nvram]}` does not reboot the attached BPEs. To reboot attached BPEs, you must use the **all** option.

Example

The following example reboots the switch at 8:00 AM on April 15, 2005:

```
reboot time 04 15 2005 08 00 00
```

History

This command was first available in ExtremeXOS 10.1.

The alternate BootROM image was added in ExtremeXOS 11.1.

The **slot**, **node-address**, **stack-topology**, and **as-standby** options were added in ExtremeXOS 12.0.

The **all** option for rebooting attached BPEs was added in ExtremeXOS 22.5.

The **rolling** option was added in ExtremeXOS 31.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

refresh access-list network-zone

```
refresh access-list network-zone [zone_name | all]
```

Description

This command is used to refresh a specific network zone, or all the network zones.

Syntax Description

network-zone	Specifies the logical group of remote devices.
<i>zone_name</i>	Specifies the network_zone name.
all	Refresh all the network-zones.

Default

N/A.

Usage Guidelines

Use this command to refresh a specific network zone, or all the network zones.

When you issue the command to refresh a network-zone, or all network-zones, it can take a long time to clear the CLI because each individual policy must be converted before it is refreshed. The command succeeds, or fails, only after it receives a response for all policy refresh results from the hardware.

If the refresh fails for a specific zone, the following error message will be printed on the console.

```
Switch # refresh access-list network-zone zone1
ERROR: Refresh failed for network-zone "zone1".
```

Example

The following example refreshes all policies in “zone1”:

```
refresh access-list network-zone zone1
```

History

This command was first available in ExtremeXOS 15.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

refresh identity-management role

```
refresh identity-management role user [user_name {domain domain_name} |
all {role role_name}]
```

Description

Refreshes the role evaluation for the specified user, for all users, or for all users currently under the specified role.

Syntax Description

<i>user_name</i>	Specifies a user name for which role evaluation will be refreshed.
<i>domain_name</i>	Specifies a domain name for the specified user.
all	Specifies a refresh for all users associated with the specified role.
<i>role_name</i>	Specifies a a role name for which all users will be refreshed.

Default

N/A.

Usage Guidelines

It may be necessary to refresh the role of a user due to a new role which might be better suited for the user or due to a change in LDAP attributes of the user which in turn might result in the user being classified under a different role. This command can be used in all such cases.

Example

The following example refreshes the role for user Tony:

```
* Switch.22 # refresh identity-management role user "Tony"
```

The following example refreshes the role for all users who are currently classified under the Marketing role:

```
* Switch.22 # refresh identity-management role all "Marketing"
```

History

This command was first available in ExtremeXOS 12.7.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

refresh igmp ssm-map

```
refresh igmp ssm-map {dns group} [grpipaddress netmask | ipNetmask]  
{{vr} vrname}
```

Description

Refreshes an *IGMP* SSM mapping entry.

Syntax Description

dns group	Refreshes DNS sources for a multicast group.
<i>grpipaddress</i>	Specifies the multicast group IP address.
<i>netmask</i>	Specifies the multicast group netmask.
<i>ipNetmask</i>	Specifies the multicast group IP address and netmask.
<i>vrname</i>	Specifies the name of the virtual router.

Default

N/A

Usage Guidelines

None.

Example

The following example refreshes an IGMP SSM mapping entry.

```
refresh igmp ssm-map 224.0.0.5/24 VR-Default
```

History

This command was first available in ExtremeXOS 15.7.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

refresh mld ssm-map

```
refresh mld ssm-map { v6groupnetmask } [{vr} vr_name]
```

Description

Sends a DNS request for a particular group. On receiving the DNS response, the “DNS Age” in the SSM mapping entry is refreshed.

Syntax Description

<i>v6groupnetmask</i>	Refreshes the specific group information.
vr <i>vr_name</i>	Specifies the virtual router name.

Default

Disabled.

Usage Guidelines

Use this command to send out DNS requests for a particular group. On receiving the DNS response, the “DNS Age” in the SSM mapping entry is refreshed.

Example

The following command send out DNS requests:

```
refresh mld ssm-map
```

When *v6groupnetmask* is specified, the SSM Mapping status and the SSM Mapping entries specific to the group range on the VR are displayed.

History

This command was first available in ExtremeXOS 15.5.

Platform Availability

This command is available on the platforms listed for the IPv6 multicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

refresh policy

refresh policy *policy-name*

Description

Refreshes the specified policy.

Syntax Description

<i>policy-name</i>	Specifies the policy to refresh.
--------------------	----------------------------------

Default

N/A.

Usage Guidelines

Use this command when a new policy file for a currently active policy has been downloaded to the switch, or when the policy file for an active policy has been edited. This command reprocesses the text file and updates the policy database.

Before 12.6.1 there was no support to refresh the policies that are associated to the local VPP. For network VPP, you can achieve policy refresh by changing the policy timestamp file. Beginning in release 11.4, the policy manager uses Smart Refresh to update the ACLs. When a change is detected, only the ACL changes needed to modify the ACLs are sent to the hardware, and the unchanged entries remain. This behavior avoids having to blackhole packets because the ACLs have been momentarily cleared. Smart Refresh works well for minor changes, however, if the changes are too great, the refresh reverts to the earlier behavior. To take advantage of Smart Refresh, disable access-list refresh blackholing by using the command:

`disable access-list refresh blackhole`

If you attempt to refresh a policy that cannot take advantage of Smart Refresh while blackholing is enabled, you will receive a message similar to the following:

Incremental refresh is not possible given the configuration of policy <name>. Note, the current setting for Access-list Refresh Blackhole is Enabled. Would you like to perform a full refresh? (Yes/No) [No]:

If blackholing is not enabled, you will receive a message similar to the following:

Incremental refresh is not possible given the configuration of policy <name>. Note, the current setting for Access-list Refresh Blackhole is Disabled. WARNING: If a full refresh is performed, it is possible packets that should be denied may be forwarded through the switch during

the time the access list is being installed. Would you like to perform a full refresh? (Yes/No) [No]:

If you attempt to refresh a policy that is not currently active, you will receive an error message.

For an ACL policy, the command is rejected if there is a configuration error or hardware resources are not available.

Example

The following example refreshes the policy zone5:

```
refresh policy zone5
```

History

This command was first available in ExtremeXOS 11.0.

Smart Refresh was added in ExtremeXOS 11.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

reset inline-power ports

```
reset inline-power ports port_list
```

Description

Power cycles the specified ports.

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports for which power is to be reset.
------------------	--

Default

N/A.

Usage Guidelines

This command power cycles the specified ports. Ports are immediately disabled and then re-enabled, allowing remote PDs to be power-cycled.

This command affects only inline power; it does not affect network connectivity for the port(s).

Example

The following command resets power for port 4 on a switch:

```
reset inline-power ports 4
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on the *PoE* devices listed in [Extreme Networks PoE Devices](#).

restart ports

```
restart ports [all | port_list]
```

Description

Resets autonegotiation for one or more ports by resetting the physical link.

Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

N/A.

Usage Guidelines

N/A.

Example

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 ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

restart process

```
restart process [class cname | name {msm slot}]
```

Description

Terminates and restarts the specified process during a software upgrade on the switch.

Syntax Description

<i>cname</i>	Specifies the name of the process to restart. With this parameter, you can terminate and restart all instances of the process associated with a specific routing protocol on all VRs. You can restart the <i>OSPF</i> routing protocol and associated processes.
<i>name</i>	Specifies the name of the process to terminate and restart. You can use this command with the following processes: bgp, eaps, exsshd, isis, lldp, netLogin, netTools, ntp, ospf, ospfv3, snmp Subagent, snmpMaster, telnetd, tftpd, vrrp, and xmld.
<i>slot</i>	On a SummitStack, specifies the node's slot number. The number is a value from 1 to 8.

Default

N/A.

Usage Guidelines

Use this command to terminate and restart a process during a software upgrade on the switch. You have the following options:

- *cname*—Specifies that the software terminates and restarts all instances of the process associated with a specific routing protocol on all VRs.
- *name*—Specifies the name of the process.

Depending on the software version running on your switch and the type of switch you have, you can terminate and restart different or additional processes. To see which processes you can restart during a software upgrade, enter `restart process` followed by `[Tab]`. The switch displays a list of available processes.

SummitStack Only

You can issue this command only from the master node. If you issue this command from any other node, the following message appears:

```
Error: Processes created by user can only be restarted on the primary node slot.
```

To display the status of ExtremeXOS processes on the switch, including how many times a process has been restarted, use the `show process {name} {detail} {description} {slotslotid}` command. The following is a truncated sample of the show process command on a switch:

Process Name	Version	Restart	State	Start Time
aaa	3.0.0.2	0	Ready	Thu Sep 1 17:00:52 2005
acl	3.0.0.2	0	Ready	Thu Sep 1 17:00:54 2005
bgp	Not Started	0	No license	Not Started
cfgmgr	3.0.0.21	0	Ready	Thu Sep 1 17:00:52 2005
cli	3.0.0.22	0	Ready	Thu Sep 1 17:00:52 2005
devmgr	3.0.0.2	0	Ready	Thu Sep 1 17:00:52 2005
dirser	3.0.0.2	0	Ready	Thu Sep 1 17:00:51 2005
dosprotect	3.0.0.1	0	Ready	Thu Sep 1 17:00:56 2005
eaps	3.0.0.8	0	Ready	Thu Sep 1 17:00:53 2005
...				

You can also use the `restart process` command when upgrading a software modular package. For more information, see the section *Upgrading a Modular Software Package* in the [Switch Engine 32.3 User Guide](#).

Example

The following example stops and restarts the process `tftpd` during a software upgrade:

```
restart process tftpd
```

The following example stops and restarts all instances of the OSPF routing protocol for all VRs during a software upgrade:

```
restart process class ospf
```

History

This command was first available in ExtremeXOS 11.3.

Support for restarting the Link Layer Discovery Protocol (lldp), Open Shortest Path First (ospf), and network login (netLogin) processes was added in ExtremeXOS 11.3.

Support for Border Gateway Protocol (bgp) and Ethernet Automatic Protection Switching (eaps) was added in ExtremeXOS 11.4.

Support for MultiProtocol Label Switching (mpls) and Virtual Router Redundancy Protocol (vrrp) was added in ExtremeXOS 11.6.

Support for netTools was added in ExtremeXOS 12.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

restart process mpls

restart process mpls

Description

Restarts the *MPLS* process when it does not respond to the CLI commands.

Default

N/A.

Usage Guidelines

None.

Example

The following command restarts the MPLS process:

```
restart process mpls
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

restart vm

restart **vm** *vm_name* {**forceful** | **graceful**}

Description

Restarts (reboots) a virtual machine (VM).

Syntax Description

vm	Virtual machine.
<i>vm_name</i>	Specifies the VM to restart.
forceful	Forcefully terminates the VM.
graceful	Gracefully shuts down the VM if possible (default).

Default

By default, the VM is gracefully shut down before restarting, if possible.

Usage Guidelines

N/A.

Example

The following example restarts the VM "testvm" gracefully:

```
restart vm testvm gracefully
```

History

This command was first available in ExtremeXOS 30.4.

Platform Availability

This command is available on all platforms that support the Integrated Application Hosting (IAH) feature and have a Core or Premier license installed. For a list of platforms that support the IAH feature and for information about licenses, see the [Switch Engine 32.3 Feature License Requirements](#).

resume vm

```
resume vm vm_name
```

Description

Resumes a virtual machine (VM) that has been suspended.

Syntax Description

vm	Specifies a VM.
<i>vm_name</i>	Specifies the VM name to resume.

Default

N/A.

Usage Guidelines

When you resume a VM, the CPU state of the VM continues from the point at which it was suspended.

To suspend the VM again, use the suspend **vm** *vm_name* command.

Example

The following example resumes the VM "vm1":

```
# resume vm vm1
```

History

This command was first available in ExtremeXOS 30.7.

Platform Availability

This command is available on all platforms that support the Integrated Application Hosting (IAH) feature and have a Core or Premier license installed. For a list of platforms that support the IAH feature and for information about licenses, see the [Switch Engine 32.3 Feature License Requirements](#).

return

```
return statusCode
```



Note

This is a script command and operates only in scripts or on the command line when scripting is enabled with the following command: `enable cli scripting {permanent}`.

Description

Exits the current script and sets the \$STATUS variable.

Syntax Description

<code>statusCode</code>	Specifies a integer value to which the \$STATUS variable is set.
-------------------------	--

Default

N/A.

Usage Guidelines

When used in nested scripts, this command allows you to terminate the current script, set the \$STATUS variable, return to the parent script, and evaluate the \$STATUS variable in the parent script. For more information on the \$STATUS variable, see “Using CLI Scripting” in the [Switch Engine 32.3 User Guide](#).

Example

The following example exits the current script and sets the \$STATUS variable to -200:

```
return -200
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

rm

```
rm file_name
```

Description

Deletes an existing configuration, policy, or if configured, core dump file from the system.

Syntax Description

<i>file_name</i>	Specifies the name of the configuration, policy file, or if configured, the core dump file to delete.
------------------	---

Default

N/A.

Usage Guidelines

After you delete a configuration or policy file from the system, that file is unavailable to the system.

You cannot delete an active configuration file (the configuration currently selected to boot the switch). To see which configuration that you are currently using, run the `show switch {detail}` command. If you attempt to delete the active configuration file, the switch displays a message similar to the following:

```
Error: Cannot remove current selected active configuration file.
```

When you delete a file from the switch, a message similar to the following appears:

```
Remove testpolicy.pol from /usr/local/cfg? (y/N)
```

Type `y` to delete the file from your system. Type `n` to cancel the process and keep the file on your system.

Case-Sensitive File Names

File names are case-sensitive. For example, if you have a configuration file named `Test.cfg`, and you attempt to delete a file with the incorrect case, for example `test.cfg`, the system is unable to delete the file. The switch does not display an error message; however, the `ls` command continues to display the file `Test.cfg`. To delete the file, make sure you use the appropriate case.

Local File Name Character Restrictions

When specifying a local file name, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z).
- Numerals (0-9).
- Period (.).
- Dash (-).
- Underscore (_).

Core Dump Files

When you configure the switch to send core dump (debug) information to internal memory, specify the file path `/usr/local/tmp`.

For information about configuring and saving core dump information, see the `configure debug core-dumps [off | directory_path]` and `save debug tracefiles directory_path` commands.

You can use the `*` wildcard to mass delete core dump files. Currently running and in-use files are not deleted.

If you configure the switch to save core dump files to internal memory and attempt to download a new software image, you might have insufficient space to complete the image download. When this occurs, you must decide whether to continue the software download or move or delete the core dump files from internal memory. To resolve this problem, if you have a switch with a USB storage device installed with space available, transfer the files to the USB device. Another option is to transfer the files from internal memory to a TFTP server. This frees up space on the internal memory while keeping the core dump files.

Example

The following example deletes the configuration file named `Activeb91.cfg` from the system:

```
# rm Activeb91.cfg
```

The following example deletes all of the core dump files stored in internal memory:

```
# rm /usr/local/tmp/*
```

The following example deletes the policy file named `test.pol` from the USB storage device:

```
# rm /usr/local/ext/test.pol
```

The following example deletes all of the configuration files from the USB storage device:

```
# rm /usr/local/ext/*.cfg
```

History

This command was first available in ExtremeXOS 10.1.

The **memorycard** option was added in ExtremeXOS 11.1.

The **internal-memory** option was added in ExtremeXOS 11.4.

Support for USB 2.0 storage devices was added in ExtremeXOS 12.5.3.

Path name support was added in ExtremeXOS 15.5.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

rmmdir

```
rmmdir directory_name
```

Description

Removes an existing directory from the specified file system or relative to the current working directory.

Syntax Description

rmmdir	Change current working directory.
<i>directory_name</i>	Pathname of a directory.

Default

N/A.

Usage Guidelines

Use this command to remove an existing directory from the specified file system or relative to the current working directory.

History

This command was first available in ExtremeXOS 15.5

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

rtlookup rpf

```
rtlookup [ipaddress | ipv6address] rpf {vr vr_name}
```

Description

Displays the RPF for a specified multicast source.

Syntax Description

<i>ipaddress</i>	Specifies an IPv4 address.
<i>ipv6address</i>	Specifies an IPv6 address.
rpf	Selects the RPF for the specified multicast source.
<i>vr_name</i>	Specifies the VR or VRF for which to display the route.

Default

vr_name is the VR of the current CLI context.

Usage Guidelines

None.

Example

The following example displays the RPF lookup for a multicast source through *VR-Default*:

```
rtlookup 2001db8::ef80:2525:1023:5213 rpf vr vr-default
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

rtlookup

```
rtlookup [ipaddress | ipv6address] { unicast | multicast | vr vr_name}
```

Description

Displays the available routes to the specified IPv6 address.

Syntax Description

<i>ipaddress</i>	Specifies an IPv4 address.
<i>ipv6address</i>	Specifies an IPv6 address.
unicast	Displays the routes from the unicast routing table in the current router context.

multicast	Displays the routes from the multicast routing table in the current router context.
<i>vr_name</i>	Specifies the VR or VRF for which to display the route.

Default

N/A.

Usage Guidelines

None.

Example

The following command performs a look up in the route table to determine the best way to reach the specified IPv6 address:

```
rtlookup 2001:db8::ef80:2525:1023:5213 unicast
```

History

This command was first available in ExtremeXOS 10.1.

The xhostname option was removed in ExtremeXOS 11.0.

Support for IPv6 was added in ExtremeXOS 11.2.

The unicast and multicast options were added in ExtremeXOS 12.1.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

run diagnostics

```
run diagnostics [extended | normal] }
```

Description

Runs normal or extended diagnostics on the switch or node, and stacking ports.

This command is not supported in stacking mode, but if you issue the show diagnostics command from the master node, it will show the diagnostic results for all the nodes.

Syntax Description

extended	Runs an extended diagnostic routine. Takes the ports offline, and performs extensive ASIC and packet loopback tests on all of the ports.
normal	Runs a normal diagnostic routine. Takes the ports offline, and performs a simple ASIC and packet loopback test on all of the ports.

Default

N/A.

Usage Guidelines

Depending on your platform, use this command to run diagnostics on the switch or stack port.

If you run the diagnostic routine on the switch, it reboots and then performs the diagnostic test. During the test, traffic to and from the ports on the switch is temporarily unavailable. When the diagnostic test is complete, the switch reboots and becomes operational again.

To run the diagnostic routine on the stack ports, you need a dedicated stacking cable that connects stack port 1 to stack port 2, which are located at the rear of the switch. The stacking cable is available from Extreme Networks. The switch performs a hardware test to confirm that the stack ports are operational; traffic to and from the ports on the switch is temporarily unavailable. This Bit Error Rate Test (BERT) provides an analysis of the number of bits transmitted in error.

After the switch runs the diagnostic routine, test results are saved to the switch's EEPROM and messages are logged to the syslog.

To run diagnostics on a switch that is in a SummitStack, first disable stacking on that switch, then restart the switch. Once restarted, log into the switch via its console port, and run diagnostics. The switch will perform the diagnostic tests, and then restart. Once restarted, log into the switch via its console port and enable stacking, then reboot the switch. Once restarted, the switch will rejoin the stack.

Viewing Diagnostics

To view results of the last diagnostics test run, use the following command:

```
show diagnostics {slot [slot_number]}
```

If the results indicate that the diagnostic failed on a node, replace the node with another switch of the same type.

If the results indicate that the diagnostic failed on the switch, contact Extreme Networks Technical Support.

The following example runs normal diagnostics on a ExtremeSwitching series switch:

```
run diagnostics normal
```

The switch displays a warning similar to the following about the impact of this test. You also have the opportunity to continue or cancel the test:

```
Running Diagnostics will disrupt network traffic.
Are you sure you want to continue? (y/n)
```

Enter **y** to continue and run the diagnostics. Enter **n** to cancel the operation.

The following command runs diagnostics on the stack ports on a ExtremeSwitching series switch:

```
run diagnostics stack-port
```

If you issue this command with a console connection, the switch displays the following information. You also have the opportunity to continue or cancel the test:

```
Summit Diagnostics Mode Enabled, Starting Diagnostics....
Found X440a-24T in Motherboard
Motherboard CPLD Revision: 2
Starting stacking port diagnostics
*****
*                                                                 *
* Please connect a cable between Stack Port 1 and Stack Port 2. *
*                                                                 *
*           Press S to skip test, ENTER key to continue.         *
*                                                                 *
*****
```

Press **[Enter]** to continue and run the diagnostics. Enter **s** to cancel the operation.

If you continue with diagnostics, the switch displays messages similar to the following:

```
Stack Port 1 and Stack Port 2
BERT .....
.
.....
Stacking ports
Port 1 (Device 0 - Device port 26)
Lane 0 PASSED.
Lane 1 PASSED.
Lane 2 PASSED.
Lane 3 PASSED.
Port 2 (Device 0 - Device port 27)
Lane 0 PASSED.
Lane 1 PASSED.
Lane 2 PASSED.
Lane 3 PASSED.
DIAGNOSTIC PASS: run test bert stacking
Summit Diagnostics completed, rebooting system...
```

If you issue this command with a Telnet connection, the switch displays a warning similar to the following about the impact of this test. You also have the opportunity to continue or cancel the test:

```
Running Diagnostics will disrupt network traffic.
Are you sure you want to continue? (y/n)
```

Enter **y** to continue and run the diagnostics. Enter **n** to cancel the operation.

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

run elrp

```
run elrp {vlan} vlan_name {ports [ports | all | none]} {remote-endpoints
vxlan all} {interval interval {seconds | milliseconds} } {retry
count}
```

Description

Starts one-time, non-periodic ELRP packet transmission on the specified ports of the VLAN using the specified count and interval.

Syntax Description

vlan <i>vlan_name</i>	Specifies a VLAN name.
<i>ports</i>	Specifies the set of VLAN ports for packet transmission.
remote-endpoints	Specifies remote endpoints that are part of this VLAN.
vxlan	Specifies VXLAN remote endpoints that are part of this VLAN.
interval	Time interval between two successive ELRP PDUs.
<i>interval</i>	Interval value between 1–64 seconds or 100–64,000 milliseconds. Default is 1 second.
seconds	Specifies that time interval is in the unit of seconds.
milliseconds	Specifies that time interval is in the unit of milliseconds.
<i>count</i>	Specifies the number of times ELRP packets must be transmitted. The range is 3 to 255 times. The default is 10 times.

Default

Second—The interval between consecutive packet transmissions is 1 second.

Count—The number of time ELRP packets must be transmitted is 10.

If ports are not specified, the command is applied to all ports.

Usage Guidelines

This command starts one-time, non-periodic ELRP packet transmission on the specified ports of the VLAN using the specified count and interval. If any of these transmitted packets is returned, indicating

loopback detection, the ELRP client prints a log message to the console. There is no need to send a trap to the *SNMP* manager for non-periodic requests.

If you do not specify the optional interval or retry parameters, the default values are used.

Use the `configure elrp-client periodic` command to configure periodic transmission of ELRP packets.

The ELRP client must be enabled globally in order for it to work on any VLANs. Use the `enable elrp-client` command to globally enable the ELRP client.

The ELRP client can be disabled globally so that none of the ELRP VLAN configurations take effect. Use the `disable elrp-client` command to globally disable the ELRP client.

Example

The following command starts one-time, non-periodic ELRP packet transmission on the VLAN green using the default interval and packet transmission:

```
run elrp green
```

History

This command was first available in ExtremeXOS 11.1.

The ability to specify the time interval in milliseconds was introduced in ExtremeXOS 22.4.

VXLAN remote endpoint option added in ExtremeXOS 22.4.

Designating ports made optional in ExtremeXOS 22.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

run failover

```
run failover {force}
```

Description

Causes a user-specified node failover.

Syntax Description

force	Force failover to occur.
--------------	--------------------------

Default

N/A.

Usage Guidelines

Use this command to cause the master node to failover to the backup node in SummitStack.

Before you initiate failover, use the `show switch {detail}` command to confirm that the nodes are in sync and have identical software and switch configurations. If the output shows MASTER and BACKUP (InSync), the two nodes are in sync.

If the master and backup SummitStack nodes' software and configuration are not in sync and are running ExtremeXOS 12.0 or later, use the `synchronize` command to get the two nodes in sync. This command ensures that the backup has the same software in flash as the master.



Note

Both the backup and the master nodes must be running ExtremeXOS 11.0 or later to use the `synchronize` command.

Example

The following command causes a failover on a SummitStack:

```
run failover
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available only on a SummitStack.

run script

```
run script filename {arg1} {arg2} ... {arg9}
```

Description

Run (plays back) an ASCII-formatted configuration file or a user-written script file on the switch. This command is synonymous with the `load script` command.

Syntax Description

<i>filename</i>	Specifies the user-defined name of the ASCII-formatted configuration file or a user-written script file. The script file is known as the XOS script file and uses the .xsf file extension.
arg	Specifies up to nine variable values that can be specified by the user. The variables are created with the names CLI.ARGV1, CLI.ARGV2, ... CLI.ARGV9.

Default

N/A.

Usage Guidelines

Use this command to load an ASCII-formatted configuration file or a user-written script file.

Configuration File: After downloading the configuration file from the TFTP server, this command loads and restores the ASCII-formatted configuration file to the switch.

An ASCII-formatted configuration file uses the .xsf file extension, not the .cfg file extension. The .xsf file extension (known as the XOS script file) saves the XML-based configuration in an ASCII format readable by a text editor.

For more detailed information about the ASCII configuration file, including the steps involved to upload, download, and save the configuration, see the `upload configuration [hostname | ipaddress] filename {vr vr-name}` command.

User-Written Script File: After writing a script, this command executes the script and passes arguments to it. As with the configuration files, these files use the .xsf file extension that is automatically added.

The command allows up to nine optional variable values to be passed to the script. These are created with the names CLI.ARGV1, CLI.ARGV2, CLI.ARGV3, ... CLI.ARGV9.

In addition, two other variables are always created. CLI.ARGV0 gives the count of the number of parameters passed, and CLI.ARGV0 contains the name of the script that is being executed.

To check the variable values use the command, `show var`.



Note

Only the .xsf extension is used. The load script command assumes an .xsf extension and retries opening the file if the file cannot be found with the original specified name or no extension is provided.

Example

The following command loads the ASCII-formatted configuration named configbackup.xsf:

```
load script configbackup.xsf
```

After issuing this command, the ASCII configuration quickly scrolls across the screen. The following is an example of the type of information displayed when loading the ASCII configuration file:

```
script.meg_upload_config1.xsf.389 # enable snmp access
script.meg_upload_config1.xsf.390 # enable snmp traps
script.meg_upload_config1.xsf.391 # configure mstp region purple
script.meg_upload_config1.xsf.392 # configure mstp revision 3
script.meg_upload_config1.xsf.393 # configure mstp format 0
script.meg_upload_config1.xsf.394 # create stpd s0
```


ExtremeXOS 15.6 provided capability for Python scripting. Current Python scripting implementation allows a script to interact directly with the CLI interface for managing ExtremeXOS functionality. Python script files end in .py. The .py suffix on the *script* file name tells the `run script` command to use the Python interpreter to process the script file. This command is functions exactly as `load script`.

History

This command was first available in ExtremeXOS 15.6.

Scripting support for Python was added in ExtremeXOS 15.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

run tech-support report

```
run tech-support report {now | in hours | cancel} {collector [all |
  hostname | ip_address]}
```

Description

This command instructs the switch to generate a report and upload it to a collector.

Syntax Description

now	Specifies that you run a report immediately. This is the default setting.
in	Specifies that you run report in a specified number of hours.
<i>hours</i>	Specifies the hours from now to run report. The range is 1-168 hours (one week).
cancel	Cancels the scheduled report.
collector	Specifies the report collector. The default value is all collectors.
all	Specifies all report collectors.
<i>hostname</i>	Specifies the host name of the collector.
<i>ip_address</i>	Specifies the IPv4 address of the collector.

Default

The default time for running reports is **now**.

The default for number of collectors is **all**.

Usage Guidelines

This command instructs the switch to generate a report and upload it to a collector. The default operation is to perform this operation immediately for all existing collectors. Optionally, you can

configure a one-time trigger to perform the operation in "hours from now." The valid range is one to 168 hours (one week). If you specify the hostname or IP address, the switch runs a report for that particular collector.

Only a single one-time report per collector can be scheduled at any time. When `run tech-support report in hours` is issued before the previous scheduled one-time report completes, the previous report is cancelled, and a new one-time report is scheduled.

This command also provides a way to cancel a scheduled report for a particular collector.

Example

The following command example configures a specific collector to display a detailed output set:

```
# run tech-support report
Connecting to 10.5.2.107:800 with SSL disabled...
Collector connected successfully.
Generating summary report for 10.5.2.107:9998.....
Report generated successfully.
Sending report to 10.5.2.107:800...
Report sent successfully

Connecting to 10.5.2.107:800 with SSL enabled...
Error: Failed to connect to the collector - Socket time out *

# run tech-support report in 1
Run tech-support report is scheduled on Thu Feb 21 05:06:32 2013 for the
collector 10.5.2.108:800.

Run tech-support report is scheduled on Thu Feb 21 05:06:32 2013 for the
collector 10.5.2.107:9998.

To cancel a scheduled report, use 'run tech-support report cancel' command.
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

run update

```
run update
```

Description

Activates a newly installed modular software package.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

After you install a modular software package to the active partition, use this command to make the update active. This command causes the ExtremeXOS system to start the newly installed processes contained in the package, without rebooting the switch.

If you installed the package to the inactive partition, you need to reboot the switch to activate the package.

Example

The following command activates any newly installed modular software packages installed on the active partition:

```
run update
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

run upm profile

```
run upm profile profile-name {event event-name} {variables variable-string}
```

Description

Executes the specified Universal Port profile on the switch.

Syntax Description

<i>profile-name</i>	Specifies the UPM profile to be run.
<i>event-name</i>	Specifies an event type for the specified profile. Valid event types are device-detect, device-undetected, user-authenticate, and user-unauthenticated.
<i>variable-string</i>	Specifies a string of variable names and the assigned variable values to be used in the profile. The format is: <i>var_name1=value_1; var_name2=value_2; var_name3=value_3</i> . Each variable name is followed by the equal sign (=), the variable value, and a semicolon (;).

Default

N/A.

Example

The following command runs a UPM profile called example on the switch:

```
run upm profile example
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the Universal Port feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

run vm-tracking repository

```
run vm-tracking repository sync-now
```

Description

Manually starts FTP file synchronization for NVPP and VM MAP files.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Before you can manually start FTP file synchronization, you must configure FTP servers using the [configure vm-tracking repository](#) command.

Example

The following command starts file synchronization with the configured FTP server:

```
# run vm-tracking repository sync-now
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

VLAN

save configuration

```
save configuration {primary | secondary | existing-config | new-config}
```

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.
<i>existing-config</i>	Specifies an existing user-defined configuration.
<i>new-config</i>	Specifies a new user-defined configuration.

Default

Saves the current configuration to the location used on the last reboot.

Usage Guidelines

The configuration takes effect on the next reboot.

Each file name must be unique and can be up to 32 characters long but cannot include any spaces, commas, or special characters.

Configuration files have a .cfg file extension. When you enter the name of the file in the CLI, the system automatically adds the .cfg file extension. Do not use this command with ASCII-formatted configuration files. Those configuration files have an .xsf file extension. For more information about using ASCII-formatted configuration files see the [upload configuration \[hostname | ipaddress\] filename {**vr** *vr-name*}](#) and the [load script filename {**arg1**} {**arg2**} ... {**arg9**}](#) commands.

This command also displays in alphabetical order a list of available configurations. The following is sample output that displays the primary, secondary, and user-created and defined configurations (“test” and “XOS1” are the names of the user-created and defined configurations):

```
exsh.9 # save configuration
<cr>                Execute the command
primary             Primary configuration file
secondary           Secondary configuration file
<existing-config>   Existing configuration file name
```

```
"test"  "XOS1"
<new-config>      New configuration file name
```

The switch prompts you to save your configuration changes. Enter y to save the changes or n to cancel the process.

If you enter n, the switch displays a message similar to the following:

```
Save configuration cancelled.
```

If you enter y, the switch saves the configuration and displays a series of messages. The following sections provide information about the messages displayed when you save a configuration on your switch.



Note

Configuration files are forward-compatible only and not backward-compatible. That is, configuration files created in a newer release, such as ExtremeXOS 12.4, might contain commands that do not work properly in an older release, such as ExtremeXOS 12.1.

Local Filename Character Restrictions

This section provides information about the characters supported by the switch for local filenames.

When specifying a local filename, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-)
- Underscore (_)

When naming a local file, remember the requirements listed above.

Saving a New Configuration

If you create and save a configuration with a new file name, the switch saves the new configuration and then prompts you to select the newly created configuration as the switch's default configuration.

The following sample output is similar to the message displayed:

```
Do you want to save configuration to test1.cfg? (y/n) Yes
Saving configuration..... done!
Configuration saved to test1.cfg successfully.
```

The switch then prompts you to select which configuration to use to bootup the system. The following sample output is similar to the message displayed:

```
The current selected default configuration database to boot up the system
(primary.cfg) is different than the one just saved (test.cfg).
Do you want to make test.cfg the default database? (y/n)
```

Enter **y** to use the new configuration as the default configuration. Enter **n** to cancel the operation and keep using the current default, active configuration.

Saving an Existing Configuration

If you make and save changes to an existing configuration, the switch prompts you to save and override the existing configuration.

The following sample output is similar to the message displayed:

```
The configuration file test.cfg already exists.  
Do you want to save configuration to test.cfg and overwrite it? (y/n) Yes  
Saving configuration ..... done!  
Configuration saved to test.cfg successfully.
```

The following sample output on a SummitStack is similar to the output displayed:

```
The configuration file primary.cfg already exists.  
Do you want to save configuration to primary.cfg and overwrite it? (y/N) Yes  
Saving configuration on primary ..... done!  
Synchronizing configuration to backup .... done!  
Saving config on Standbys (Slots: 1).  
...  
Configuration saved on Standby (Slot 1): done!
```

If you override an existing configuration that is not the current default, active configuration, the switch prompts you to select which configuration to use to bootup the system. The following sample output is similar to the message displayed:

```
The current selected default configuration database to boot up the system  
(primary.cfg) is different than the one just saved (test.cfg).  
Do you want to make test.cfg the default database? (y/n) No  
Default configuration database selection cancelled.
```

Enter **y** to use the updated configuration as the default configuration. Enter **n** to cancel the operation and keep using the current default, active configuration.

Example

The following command saves the current switch configuration to the configuration file named **XOS1**:

```
save configuration XOS1
```

The following command save the current switch configuration to the secondary configuration file:

```
save configuration secondary
```

History

This command was first available in ExtremeXOS 10.1.

The status messages displayed by the switch were updated in ExtremeXOS 11.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

save configuration as-script

```
save configuration as-script script-name
```

Description

Saves the running configuration as a script.

Syntax Description

<i>script-name</i>	Specifies the name of the file to save the configuration to. The script file is known as the XOS script file and uses the .xsf file extension.
--------------------	--

Default

N/A.

Usage Guidelines

This command allows you to save the current configuration as a script and export it out of the box for later use.

For SummitStack only

The script is saved on all the nodes in a SummitStack when the save configuration as-script command is executed.

Example

The following example saves a running ASCII-formatted configuration named primary.xsf.

```
save configuration as-script primary.xsf
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

save configuration automatic

```
save configuration automatic {every minutes {primary | secondary |
    existing-config | new-config} | never}
```

Description

This command configures the periodic auto-save of the currently running switch configuration.

Syntax Description

automatic	Configures auto-save of system configuration.
every	Sets the switch configuration to be saved at the designated recurrent intervals.
<i>minutes</i>	Designates the auto-save interval in minutes with a range of 2-1,440 minutes (default is two minutes).
primary	Designates the primary configuration file for saving.
secondary	Designates the secondary configuration file for saving.
<i>existing-config</i>	Name of the existing configuration file name.
<i>new-config</i>	New configuration file name.
never	Turns off auto-save feature (default is turned off).

Default

By default, auto-save is turned off.

If you do not select a time interval for saving, the default is two minutes.

By default, the configuration is saved to the file specified in the Config Automatic field of the [show switch](#) on page 3243 command output. If no value appears in this field, the configuration is saved to the file specified in the Config Selected field of the [show switch](#) on page 3243 command. If no value appears in this field either, the configuration is saved to `autosave.cfg`.

Usage Guidelines

The auto-save features is turned off by default. To turn on the auto-save feature, use the command `save configuration automatic` (to accept the default two-minute save interval) or `save configuration automatic every minutes` (to specify the auto-save interval). The message `Do you want to auto-save configuration to primary.cfg and overwrite it? (y/N)` appears. Select "yes" to enable the auto-save to the `primary.cfg` file. Selecting "no" cancels the command.

To turn off auto-save, use the command `save configuration automatic never`.

If you want to specify a different file to save the configuration to (than the default `primary.cfg`), use the command `save configuration automatic {every minutes {primary |`

secondary | *existing-config* | *new-config*}}, specifying an auto-save interval and configuration file name.

To see the current status of the auto-save feature, use the command [show switch](#) on page 3243.

Example

The following example turns on auto-save, accepting the default auto-save interval (two minutes) and the default configuration file (`primary.cfg`):

```
save configuration automatic
```

The switch status appear as:

```
show switch
...
Config Selected: primary.cfg
Config Booted: primary.cfg
Config Automatic: primary.cfg

primary.cfg      Created by ExtremeXOS version 22.2.0.16
                  344404 bytes saved on Tue Jan 17 11:17:56 2017
                  Auto-saved every 2 minutes.
                  Next periodic save on Tue Jan 17 14:45:33 2017
```

The following example turns off auto-save:

```
save configuration automatic never
```

The switch status appear as:

```
show switch
...
Config Selected: primary.cfg
Config Booted: primary.cfg
Config Automatic: NONE

primary.cfg      Created by ExtremeXOS version 22.2.0.16
                  344404 bytes saved on Tue Jan 17 14:45:33 2017
```

The following example changes the auto-save interval to five minutes and makes `autosave.cfg` the file that is saved to:

```
save configuration automatic every 5 autosave
```

The switch status appears as:

```
show switch
...
Config Selected: primary.cfg
Config Booted: primary.cfg
Config Automatic: primary.cfg (Disabled)
primary.cfg Created by ExtremeXOS version 22.2.0.16
344404 bytes saved on Tue Jan 17 14:45:33 2017
Auto-save not enabled
```

History

This command was first available in ExtremeXOS 22.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

save debug tracefiles

```
save debug tracefiles directory_path
```

Description

Copies debug information to USB 2.0 storage device.

Syntax Description

<i>directory_path</i>	Directory path (memory card is /usr/local/ext; internal memory is /usr/local/tmp; and home directory is /usr/local/cfg).
-----------------------	--

Default

N/A.

Usage Guidelines



Note

Use this command only under the guidance of Extreme Networks Technical Support to troubleshoot the switch.

Use this command to copy debug information to an installed removable storage device. The debug information includes log files and trace files.

Progress messages are displayed that indicate the file being copied and when the copying is finished.

Beginning with ExtremeXOS 11.6, you can use the `upload debug [hostname | ipaddress] [{vr} vrname]` command to copy debug information to a network TFTP server.

Example

The following command copies debug information to a removable storage device:

```
# save debug tracefiles /usr/local/ext
```

History

This command was first available in ExtremeXOS 11.0.

The syntax for this command was modified in ExtremeXOS 11.1 from `upload debug-info memorycard` to `save debug tracefiles memorycard`.

Support for USB 2.0 storage devices was added in ExtremeXOS 12.5.3.

The option **memorycard** was removed and the variable *directory_path* was added in ExtremeXOS 30.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

save var key

```
save var key key [var1 var2 ...]
```



Note

This is a script command and operates only in scripts or on the command line when scripting is enabled with the following command: [enable cli scripting {permanent}](#).

Description

Saves the specified variables to the specified key.

Syntax Description

<i>key</i>	Specifies the key to which the specified variables are saved.
<i>var1 var2</i>	Specifies the variables to save, The first variable is mandatory, up to four more optional variables can be specified.

Default

N/A.

Usage Guidelines

The variables saved by the SAVE VAR command are represented by the specified key and can be retrieved and restored in the context in which this profile was applied. They are available to rollback events like user-unauthenticate and device-undetected. The key option allows the user to save data for a unique key and retrieve the saved data based on this key. The user is responsible for generating unique keys for each variable. The system has a limited amount of memory to store these variables.

Example

The following example saves the variables “username,” “ipaddr,” and “vlan” to the key “blue:”

```
save var key blue username ipaddr vlan
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

save vm image

```
save vm vm_name image image_file
```

Description

Exports a disk image of an existing virtual machine (VM).

Syntax Description

vm	Virtual machine.
<i>vm_name</i>	Specifies the VM to export a disk from.
image	Saves (exports) the disk image of a VM in current format (QCOW2 or VMDK).
<i>image_file</i>	Specifies the file name for the exported VM disk image. File extension are appended if not specified.

Default

N/A.

Usage Guidelines

N/A.

Example

The following example exports the disk image from VM "testvm" to a file named "testvmimage":

```
# save vm testvm image testvmimage
```

History

This command was first available in ExtremeXOS 30.4.

Platform Availability

This command is available on all platforms that support the Integrated Application Hosting (IAH) feature and have a Core or Premier license installed. For a list of platforms that support the IAH feature and for information about licenses, see the [Switch Engine 32.3 Feature License Requirements](#).

save vm state

```
save vm vm_name state
```

Description

Pauses the CPU state of a guest virtual machine (VM), saves the state to a file, and then moves the VM to "shut-off" state to free resources.

Syntax Description

vm	Specifies a virtual machine.
<i>vm_name</i>	Specifies the virtual machine name.
state	Specifies saving the running CPU state to persist across switch reboots.

Default

N/A.

Usage Guidelines

The saved state persists across switch reboots and virtMgr process restarts. The next time the VM is started, either by explicit `start vm vm_name` command or by autostart (enable `vm vm_name autostart`), the saved CPU state is restored.

A suspended VM (`suspend vm vm_name`) can be saved. In this case, it is returned to a suspended state after the next time it is started, and requires an explicit resume (`resume vm vm_name`) to continue processing.

You cannot apply this command to a VM with Integrated Application Hosting (IAH) dedicated ports attached.

Example

The following example saves the vm "vm1" preserving the state:

```
# save vm vm1 state
```

History

This command was first available in ExtremeXOS 30.7.

Platform Availability

This command is available on all platforms that support the Integrated Application Hosting (IAH) feature and have a Core or Premier license installed. For a list of platforms that support the IAH feature and for information about licenses, see the [Switch Engine 32.3 Feature License Requirements](#).

scp2

```
scp2 {cipher cipher} {mac mac} {compression [on | off]} {port portnum}
    {vr vr_name} user [hostname | ipaddress]:remote_file local_file
```

or

```
scp2 {cipher cipher} {mac mac} {compression [on | off]} {port portnum}
    {vr vr_name} local_file user [hostname | ipaddress]:remote_file
```

Description

The first command initiates an SCP2 client session to a remote SCP2 server and copies a configuration or policy file from the remote system to the switch.

The second command initiates an SCP2 client session to a remote SCP2 server and copies a configuration or policy file from the switch to a remote system.

Syntax Description

<i>vr_name</i>	Specifies the virtual router. The default virtual router is <i>VR-Mgmt.</i> Note: User-created VRs are supported only on the platforms listed for this feature in the Switch Engine 32.3 Feature License Requirements document.
<i>cipher</i>	Specifies the name of the cipher. Possible values are: <ul style="list-style-type: none"> 3des-cbc aes128-cbc aes128-ctr aes192-cbc aes192-ctr aes256-cbc aes256-ctr chacha20-poly1305@openssh.com rijndael-cbc@lysator.liu.se
<i>mac</i>	Specifies the name of the Message Authentication Code. Possible values are: <ul style="list-style-type: none"> hmac-md5 hmac-md5-96 hmac-md5-96-etm@openssh.com hmac-md5-etm@openssh.com hmac-sha1 hmac-sha1-96 hmac-sha1-96-etm@openssh.com hmac-sha1-etm@openssh.com hmac-sha2-256 hmac-sha2-256-etm@openssh.com hmac-sha2-512 hmac-sha2-512-etm@openssh.com

on	Specifies that the data is to be compressed.
off	Specifies that compression is not to be used. This is the default.
<i>portnum</i>	Specifies the TCP port number to be used for communicating with the SSH2 client. The default is port 22.
<i>user</i>	Specifies a login name for the remote host.
<i>hostname</i>	Specifies the name of the remote host.
<i>ipaddress</i>	Specifies the IP address of the remote host. Note: For IPv6 addresses, use square brackets.
<i>remote_file</i>	Specifies the name of the remote file (configuration file, policy file, image file, public key file) to be transferred.
<i>local_file</i>	Specifies the name of the local file (configuration file, policy file, image file, public key file) to be transferred.

Default

The default settings for SSH2 parameters are as follows:

- cipher—the full cipher list
- mac—the full Message Authentication Code list
- port—22
- compression—off
- vr_name—VR-Mgmt

Usage Guidelines

SSH2 does not need to be enabled on the switch in order to use this command.

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.

Host Name, User Name, and Remote IP Address Character Restrictions

This section provides information about the characters supported by the switch for host names and remote IP addresses.

When specifying a host name, user name, or remote IP address, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-) Permitted for host and user names
- Underscore (_) Permitted for host and user names
- Colon (:)

- At symbol (@). Permitted only for user names
- Slash (/). Permitted only for user names

When naming the host, creating a user name, or configuring the IP address, remember the requirements listed above.

Remote Filename Character Restrictions

This section provides information about the characters supported by the switch for remote filenames.

When specifying a remote filename, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-)
- Underscore (_)
- Slash (/)

When naming a remote file, remember the requirements listed above.

Example

The following command copies the configuration file test.cfg on host system1 to the switch:

```
# scp2 admin@system1:test.cfg localtest.cfg
```

The following command copies the configuration file engineering.cfg from the switch to host system1:

```
# scp2 engineering.cfg admin@system1:engineering.cfg
```

The following command copies the file Anna5.xsf from the default virtual router to 150.132.82.140:

```
# scp2 vr vr-default Anna5.xsf root@150.132.82.140:Anna5.xsf
Upload /config/Anna5.xsf to
Connecting to 150.132.82.140...
```

History

This command was first available in ExtremeXOS 11.2.

Changes to **cipher**, as well as the addition of **mac** and **compression**, were first available in ExtremeXOS 15.7.1.

Ciphers and MACs that are unsupported in OpenSSH 8.1p1 were removed in ExtremeXOS 30.7.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

set var

```
set var varname _expression
```



Note

This is a script command and operates only in scripts or on the command line when scripting is enabled with the following command: [enable cli scripting {permanent}](#).

Description

Creates and sets the CLI scripting variable to the desired value.

Syntax Description

<i>varname</i>	Specifies the name of the CLI scripting variable. Valid format is \$VARNAME (case insensitive, character string up to 32 characters).
<i>_expression</i>	Specifies the _expression whose value should be evaluated and used to set the variable.

Default

N/A.

Usage Guidelines

The format of a local variable (case insensitive) is: \$VARNAME.

An error message is displayed if the user attempts to use a variable name with a length greater than 32 characters.

If a variable already exists, it is overwritten. No error message is displayed.

Example

The following examples show some ways you can manipulate variables:

```
Set var x 100
Set var x ($x + 2)
Set var y ($x - 100)
Set var y ($ (x) - 100)
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show access-list

```
show access-list {any | ports port_list | vlan vlan_name} {ingress | egress}
```

Description

Displays the ACLs configured on an interface.

Syntax Description

<i>aclname</i>	Specifies the ACL name. The name can be from 1-32 characters long.
any	Specifies the wildcard ACL.
<i>port_list</i>	Specifies which ports' ACLs to display.
<i>vlan_name</i>	Specifies which <u>VLAN</u> 's ACL to display.
ingress	Display ingress ACLs.
egress	Display egress ACLs.

Default

The default is to display all interfaces, ingress.

Usage Guidelines

The ACL with the port and VLAN displayed as an asterisk (*) is the wildcard ACL.

If you do not specify an interface, the policy names for all the interfaces are displayed, except that dynamic ACL rule names are not displayed. To display dynamic ACLs use the following commands:

```
show access-list dynamic
```

```
show access-list dynamic rule rule {detail}
```

If you specify an interface, all the policy entries, and dynamic policy entries are displayed.

Example

The following command displays all the interfaces configured with an ACL:

```
show access-list
```

The output from this command is similar to:

Vlan Name	Port	Policy Name	Dir	Rules	Dyn Rules
=====					
*	3:6	TCP_flag	ingress	3	2
*	3:8	qos_hongkong	ingress	3	0
*	2:1	tc_2.4	ingress	4	0

*	2:7	tcp	ingress	1	0
v1	*	tcp	ingress	1	0
*	*	firewall1	ingress	2	1

The following command displays the ingress access list entries configured on the VLAN v1006:

```
show access-list v1006 ingress
```

The output from this command is similar to the following:

```
# RuleNo 1
entry dacl13 {          #Dynamic Entry
if match all {
ethernet-destination-address 00:01:05:00:00:00 ;
} then {
count c13 ;
redirect 1.1.5.100 ;
} }
# RuleNo 2
entry dacl14 {          #Dynamic Entry
if match all {
ethernet-source-address 00:01:05:00:00:00 ;
} then {
count c14 ;
qosprofile qp7 ;
} }
# RuleNo 3
entry dacl13 {
if match all {
ethernet-destination-address 00:01:05:00:00:00 ;
} then {
count c13 ;
redirect 1.1.5.100 ;
} }
```

History

This command was first available in ExtremeXOS 10.1.

The *aclname* option was removed in ExtremeXOS 11.1.

The ingress, egress, any, ports, and vlan options were added in ExtremeXOS 11.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show access-list configuration

```
show access-list configuration
```

Description

Displays the *ACL* configuration.

Syntax Description

There are no arguments or variables for this command.

Default

N/A.

Usage Guidelines

This command displays the state of the ACL configuration, set by the following commands:

```
enable access-list refresh blackhole
enable access-list permit to-cpu
configure access-list rule-compression port-counters
configure access-list vlan-acl-precedence
```

Example

The following command displays the state of the ACL configuration:

```
show access-list configuration
```

On a series switches, the output from this command is similar to the following:

```
Access-list Refresh Blackhole: Enabled
Access-list Permit To-CPU: Enabled
Access-list configured vlan-acl precedence mode: Dedicated or Shared
Access-list operational vlan-acl-precedence mode: Dedicated or Shared
Access-list Rule-compression Port-counters: Dedicated or Shared
```

The following displays how the output looks when "multiple matches" action resolution mode is chosen:

```
Access-list Refresh Blackhole: Enabled
Access-list configured vlan-acl-precedence mode: Dedicated
Access-list operational vlan-acl-precedence mode: Dedicated
Access-list Rule-compression Port-counters: Dedicated
Access-list Action Resolution: Multiple
```

The following displays how the output looks when "highest priority only" action resolution mode is chosen:

```
Access-list Refresh Blackhole: Enabled
Access-list configured vlan-acl-precedence mode: Dedicated
Access-list operational vlan-acl-precedence mode: Dedicated
Access-list Rule-compression Port-counters: Dedicated
Access-list Action Resolution: Highest Priority Only
```

The command `show configuration acl` also shows the `configure access-list action-resolution highest-priority` command if "highest priority only" action resolution mode is chosen:

```
show config acl
#
# Module acl configuration.
#
configure access-list action-resolution highest-priority
```

History

This command was first available in ExtremeXOS 11.0.

The Access-list Permit to CPU configuration was added in ExtremeXOS 11.3.2.

The Access-list Rule-compression Port Counters configuration was added in ExtremeXOS 12.3.

The Access-list Configured VLAN-ACL Precedence Mode configuration was added in ExtremeXOS 12.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show access-list counter

```
show access-list counter {countername} {any | ports port_list | vlan
    vlan_name} {ingress | egress}
```

Description

Displays the specified access list counters.

Syntax Description

<i>countername</i>	Specifies the <u>ACL</u> counter to display.
<i>port_list</i>	Specifies to display the counters on these ports.
<i>vlan_name</i>	Specifies to display the counters on the <u>VLAN</u> .
ingress	Specifies to display ingress counters.
egress	Specifies to display egress counters.

Default

The default direction is ingress.

Usage Guidelines

Use this command to display the ACL counters.

Example

The following example displays all the counters for all ACLs:

```
# sh access-list counter
Policy Name      Vlan Name      Port   Direction
Counter Name    Packet Count    Byte Count
=====
temp             *              15     ingress
  mac1           0
  mac2           2720
egressdeny       *              15     egress
  mac01count     0
temp             *              20     egress
  mac1           0
  mac2           0
```

History

This command was first available in ExtremeXOS 10.1.

The egress option was first available in ExtremeXOS 11.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show access-list counters process

```
show access-list counters process [snmp | telnet | ssh2 | http]
```

Description

Displays the access-list permit and deny statistics.

Syntax Description

snmp	Specifies statistics for <i>SNMP</i> .
telnet	Specifies statistics for Telnet.
ssh2	Specifies statistics for SSH2.
http	Specifies statistics for HTTP.

Default

N/A.

Usage Guidelines

Use this command to display the access-list permit and deny statistics. The permit and deny counters are updated automatically regardless of whether the *ACL* is configured to add counters.

Example

The following command displays permit and deny statistics for the SNMP application:

```
# sh access-list counter process snmp
```

Following is sample output for this command:

```
show access-list counter process snmp
=====
Access-list      Permit Packets      Deny Packets
=====
a1                10                0
a3                0                25
a2                0                6
=====
Total Rules : 3
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show access-list dynamic rule

```
show access-list dynamic rule [rule | rule_li ] detail
```

Description

Displays the syntax of a dynamic ACL.

Syntax Description

<i>rule</i>	Specifies the rule to display.
<i>rule_li</i>	Specifies the dynamic rule name for Lawful Intercept account only. You must have lawful intercept user privileges to specify this variable.
detail	Specifies to display where the ACL has been applied.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays the syntax of the dynamic ACL `udpACL`:

```
show access-list dynamic rule udpACL
```

The output of the command is similar to the following:

```
entry udpACL {
  if match all {
    source-address 10.203.134.0/24 ;
    destination-address 140.158.18.16/32 ;
    protocol udp ;
    source-port 190 ;
    destination-port 1200 - 1250 ;
  } then {
    permit ;
  } }
```

The following command displays where the dynamic ACL `udpACL` has been applied:

```
show access-list dynamic rule udpACL
```

The output of the command is similar to the following:

```
Rule udpACL has been applied to the following interfaces.
Vlan Name      Port      Direction
=====
*              1          ingress
```

The lawful intercept user can display the names of the existing dynamic ACLs and a count of how many times each is used when the following command is issued:

```
* show access-list dynamic
Dynamic Rules: ((*)- Rule is non-permanent )
(*)hclag_arp_0_4_96_51_fe_b2  Bound to 0 interfaces for application HealthCheckLAG
(*)idmgmt_def_blacklist      Bound to 0 interfaces for application IdentityManager
(*)idmgmt_def_whitelist      Bound to 0 interfaces for application IdentityManager
(*)mirror-data               Bound to 2 interfaces for application CLI
```

Use the following command to see the conditions and actions for a dynamic ACL:

```
* show access-list dynamic rule "mirror-data"
entry mirror-data {
  if match all {
    source-address 10.66.9.8/24 ;
    protocol udp ;
  } then {
    permit ;
    mirror law_mirror ;
  } }
```

History

This command was first available in ExtremeXOS 11.3.

The **detail** keyword was added in ExtremeXOS 11.4.

The `rule_li` variable was added in ExtremeXOS 15.3.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show access-list dynamic counter

```
show access-list dynamic counter [{countername} any | {countername}
    ports port_list | {countername} vlan vlan_name} {ingress | egress}
```

Description

Displays the dynamic ACL counters.

Syntax Description

<i>countername</i>	Display the counter.
any	Specifies the wildcard ACL.
<i>port_list</i>	Specifies which ports' ACLs to display.
<i>vlan_name</i>	Specifies which <u>VLAN</u> 's ACL to display.
ingress	Display ingress ACLs.
egress	Display egress ACLs.

Default

The default is to display all interfaces, ingress.

Usage Guidelines

None.

Example

The following command displays all the dynamic ACL counters:

```
# sh access-list dynamic counter
Vlan Name      Port  Direction
Counter Name          Packet Count      Byte Count
=====
*                15    ingress
  mac1Count          594
*                15    egress
  egressCount         0
```

The following command displays output with port number specified:

```
# sh access-list dynamic counter ports 15
Vlan Name      Port  Direction
Counter Name          Packet Count      Byte Count
=====
*                15    ingress
```

```

    mac1Count      757
*      15      egress
    egressCount      0

```

The following command displays output with counter name specified:

```

# sh access-list counter mac1 ports 15,20
Policy Name      Vlan Name      Port  Direction
Counter Name      Packet Count      Byte Count
=====
temp      *      15      ingress
    mac1      0
temp      *      20      egress
    mac1      0

```

History

This command was first available in ExtremeXOS 11.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show access-list dynamic

```
show access-list dynamic
```

Description

Displays the names of existing dynamic [ACLs](#) and a count of how many times each is used.

Syntax Description

There are no arguments or variables for this command.

Default

N/A.

Usage Guidelines

This command displays the names of existing dynamic ACLs, and how many times the ACL is used (bound to an interface).

To see the conditions and actions for a dynamic ACL, use the following command:

```
show access-list dynamic rule rule {detail}
```

Example

The following command displays names of all the dynamic ACLs:

```
show access-list dynamic
```

The following is sample output for this command:

```
Dynamic Rules:
Udpacl                Bound to 1 interfaces
icmp-echo              Bound to 1 interfaces
```

History

This command was first available in ExtremeXOS 11.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show access-list interface

```
show access-list {rule rule {start} } [ any | port port | vlan
    vlan_name ] {zone zone_name { appl-name appl_name {priority
    number }}} {ingress | egress} {detail}
```

Description

Displays the specified ACL zones, including their priority, applications, and the application priorities.

Syntax Description

any	Displays all zones on the specified interface.
port port	Displays all ACLs associated with the specified ports.
vlan vlan_name	Displays all ACLs associated with the specified <u>VLAN</u> .
zone_name	Specifies a zone to be displayed.
appl-name appl_name	Displays information by application within a zone.
priority number	Displays ACLs of the specified priority only, within an application area.
ingress	Displays ACLs applied to traffic in the ingress direction.
egress	Displays ACLs applied to traffic in the egress direction.
detail	Displays all ACLs applied to the specified interface.

Default

N/A.

Usage Guidelines

Use this command to display the ACL zones, applications, and priorities.

Specifying a zone will show all the ACLs installed in the particular zone. Specifying a priority within a zone will show all the ACLs installed at a particular priority within a zone.

Use the detail keyword to display all ACLs installed on a given interface.

Example

The following example displays the detailed view of the ACLs on port 1:1:

```
show access-list port 1:1 detail
```

The output of this command is similar to the following:

```
# show access-list port 1:1 detail
RuleNo      Application      Zone      Sub Zone
=====
      1      CLI          myZone      1
entry mac1 {
if match all {
ethernet-source-address 00:0c:29:e5:94:c1 ;
destination-address 192.168.11.144/32 ;
} then {
count mac1 ;
} }
      2      CLI          myZone      5
entry mac51 {
if match all {
ethernet-source-address 00:0c:29:e5:94:51 ;
} then {
count mack51;
} }
      3      CLI          myZone      5
entry mac52 {
if match all {
ethernet-source-address 00:0c:29:e5:94:52 ;
} then {
count mac52 ;
} }
```

The following example displays the detailed view of the priority 5 ACLs in the zone myzone on port 1:1:

```
# show access-list port 1:1 zone myZone priority 5 detail
RuleNo      Application      Zone      Sub Zone
=====
      2      CLI          myZone      5
entry mac51 {
if match all {
ethernet-source-address 00:0c:29:e5:94:51 ;
} then {
count mack51;
} }
      3      CLI          myZone      5
entry mac52 {
if match all {
```

```

ethernet-source-address 00:0c:29:e5:94:52 ;
} then {
count mac52 ;
} }

```

The following example displays the priority 5 ACLs in the zone myzone on port 1:1:

```

# show access-list port 1:1 zone myZone priority 5
#Dynamic Entries  ((*)- Rule is non-permanent )
RuleNo      Name      Application      Zone      Sub-Zone
1           mac51      CLI             myZone     5
2           mac52      CLI             myZone     5

```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show access-list meter

```

show access-list meter {metername} [any | ports [port_list ] | vlan
vlanname ] {ingress | egress}

```

Description

Displays the specified access list meter statistics and configurations.

Syntax Description

<i>meter-name</i>	Specifies the <u>ACL</u> meter to display.
out-of-profile	Show the meter out-of-profile status.
disabled-ports	Show the meter out-of-profile status that resulted in disable-port action.
<i>port_list</i>	Specifies the port list name to display the meters on.
<i>port_group</i>	Specifies a port group name to display the meters on.
<i>vlan_name</i>	Specifies to display the meters on the <u>VLAN</u> .
global-count	Counter of all the rules using a per-port meter.
ingress	ACLs applied to ingress.
egress	ACLs applied to egress.

Default

N/A.

Usage Guidelines

Use this command to display the ACL meters.

Example

The following example displays access list meter information for port 7:1:

```
# show access-list meter ports 1-4
=====
Policy Name      Vlan Name      Port
      Meter      Committed      Max Burst      Out-of-Profile  Out-of-Profile
                  Rate (Kbps)      Size      Action      DSCP      Packet Count
=====
(none)           *              2
  ingmeter3      3000000 Kbps    300000 Kb      - - - Dr      -      123456
  ingmeter4      4000000 pps     400000 pkt     - - - Dr      -           0
(none)           *              3
  ingmeter12      Max            123456 Kb      - - - Dr      -           0
  ingmeter3      3000000 Kbps    300000 Kb      - - - Dr      -           0
  ingmeter4      4000000 pps     400000 pkt     - T - Dr      -           0
(none)
  ingmeter12      Max            123456 Kb      L T D DrP     64      871234
  ingmeter3      3000000 Kbps    300000 Kb      - - D Dr      -           0
  ingmeter4      4000000 pps     400000 pkt     L - D Dr      -           0

Action   :  (D) Disable Port, (Dr) Drop, (DrP) Set Drop Precedence,
              (L) Log, (T) Trap
```

History

This command was available in ExtremeXOS 11.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show access-list network-zone

```
show access-list network-zone {zone_name}
```

Description

Displays the network-zones configured, the number of attributes configured, and the number of policy files that have the specified zones in it.

Syntax Description

network-zone	Specifies the logical group of remote devices.
<i>zone_name</i>	Specifies the network-zone name.

Default

N/A.

Usage Guidelines

Use this command to display detailed information about a particular network-zone, the attributes configured in the zone, and the policies bound to the zone.

Example

The following example displays network-zone statistics for all configured zones:

```
Switch # sh access-list network-zone
=====
Network Zone              No. of      No. of Policies
Entities      Bound
=====
zone1                      5             2
zone2                      3             1
zone3                      0             0
=====
Total Network Zones : 3
```

This example displays statistics for the specified zones, “zone1”, and “zone2”:

```
Switch #show access-list network-zone zone1
Network-zone      : zone1
Total Attributes  : 3
Attributes        : 10.1.1.1 / 32
10.1.1.1 / 30
10.1.1.0 / 24
No. of Policies   : 1
Policies          : test
Switch # sh access-list network-zone zone2
Network-zone      : zone2
No. of Entities   : 3
Entities          : 00:00:00:00:00:22 / ff:ff:ff:ff:ff:ff
00:00:00:00:00:23 / ff:ff:ff:ff:00:00
00:00:00:00:00:24 / ff:ff:ff:ff:ff:00
No. of Policies   : 0
```

History

This command was first available in ExtremeXOS 15.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show access-list usage acl-mask port

```
show access-list usage acl-mask port port
```


Description

Displays the number of *ACL* masks consumed by the ACLs on a particular port.

Syntax Description

<i>port</i>	Displays the usage on the specified port.
-------------	---

Default

N/A.

Usage Guidelines

Use this command to display how many masks are currently consumed on a port.

Example

The following example displays the ACL mask usage on port 1:1:

```
Switch.8 # show access-list usage acl-mask port 1:1
Used: 3   Available: 12
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show access-list usage acl-range port

```
show access-list usage acl-range port port
```

Description

Displays the number of Layer 4 port ranges consumed by the ACLs on the slices that support a particular port.

Syntax Description

<i>port</i>	Specifies to display the usage for the slices that support this port.
-------------	---

Default

N/A.

Usage Guidelines

ExtremeSwitching series switches can support a total of 16 Layer4 port ranges among the slices that support each group of 24 ports.

Use this command to display how many of these Layer4 ranges are currently consumed by the ACLs on the slices that support a particular port. The output of this command also displays which ports share the same slices as the specified port.

Example

The following example displays the Layer4 range usage on port 9:1:

```
Switch.3 # show access-list usage acl-range port 9:1
Ports 9:1-9:12, 9:25-9:36
L4 Port Ranges:  Used: 4  Available: 12
```

History

This command was first available in ExtremeXOS 11.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show access-list usage acl-rule port

```
show access-list usage acl-rule port port
```

Description

Displays the number of ACL rules consumed by the ACLs on a particular port or on the slices that support a particular port.

Syntax Description

<i>port</i>	Specifies to display the usage on this port.
-------------	--

Default

N/A.

Usage Guidelines

Use this command to display the rules used per slice, and also display the rule usage of the specified port.

Example

The following example displays the ACL rule usage on port 5:

```
Switch.3 # show access-list usage acl-rule port 5
Ports 1-12, 25-36
Total Rules:      Used: 34  Available: 990
```

In ExtremeXOS 15.5.1 and onwards, unless there is at least 1 rule in a given slice, the slice is not allocated. Since the slice is not yet allocated, a physical slice is not assigned to a virtual slice. So in this previous example, "used" displays what is used in that particular slice, and "available" shows the remaining rules in that particular used slice.

The following example displays the ACL ingress and egress rule usage on port 5:1:

```
Switch.4 # show access-list usage acl-rule port 5:1
Ports 5:1-5:48
Total Ingress/Egress Rules:
Used: 11  Available: 8181
Used: 1  Available: 1023
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show access-list usage acl-slice port

```
show access-list usage acl-slice port port
```

Description

Displays the number of [ACL slices](#) and rules consumed by the ACLs on the slices that support a particular port.

Syntax Description

<i>port</i>	Specifies to display the usage for the slices that support this port.
-------------	---

Default

N/A.

Usage Guidelines

Use this command to display how many slices and how many rules per each slice are currently consumed by the ACLs on the slices that support a particular port. This command also displays which ports share the same slices as the specified port.

Beginning with ExtremeXOS 12.5, you can reserve or allocate a slice for a specific feature such that rules for the feature does not share a slice with other components. A text string has been added at the end of the output for each slice that indicates which feature, if any, is reserving the slice. See the example below.

In ExtremeXOS 15.5.1 and onwards, unless there is at least one rule in a given slice, the slice is not allocated. Since the slice is not yet allocated, a physical slice is not assigned to a virtual slice. So "used" displays what is used in that particular slice, and "available" shows the remaining rules in that particular used slice.

Slices are allocated or reserved as follows:

- **user/other**—The slice is used by user ACLs and/or other switch features.
- **Reserved for:** *feature name*—the slice is reserved for the named feature, for instance VLAN statistics. Rules for this feature may not share a slice with other features or user ACLs.
- **system**—The slice contains only rules used for certain specific switch features. User ACLs may not share a slice with a system slice.

Example

The following example displays the ACL slice usage on port 1 on an ExtremeSwitching X870 series switch:

```
# show access-list usage acl-slice port 1
Ports 1-93, 95, 97, 101, 105, 109, 113, 117, 121, 125
Stage: INGRESS Pipe 0
Group 3 Priority 31 Rules:    Used:      10 Available   1014 system Double Reserved=FALSE
Reservations:
type          num          mode
Stage: INGRESS Pipe 1
Group 6 Priority 31 Rules:    Used:      0 Available   2048 user/other IntraSliceDouble
Reserved=FALSE
Group 7 Priority 30 Rules:    Used:      1 Available   2047 user/other IntraSliceDouble
Reserved=FALSE
Group 4 Priority 29 Rules:    Used:      1 Available   2047 user/other IntraSliceDouble
Reserved=FALSE
Group 5 Priority 28 Rules:    Used:     10 Available    758 system Double Reserved=FALSE
Reservations:
type          num          mode
Stage: INGRESS Pipe 2
Group 1 Priority 31 Rules:    Used:     10 Available   1014 system Double Reserved=FALSE
Reservations:
type          num          mode
Stage: INGRESS Pipe 3
Group 2 Priority 31 Rules:    Used:     10 Available   1014 system Double Reserved=FALSE
Reservations:
type          num          mode
Stage: EGRESS
Slices:
Used: 0 Available: 4
Virtual Slice * (physical slice 0) Rules:    Used:      0 Available:   256
Virtual Slice * (physical slice 1) Rules:    Used:      0 Available:   256
Virtual Slice * (physical slice 2) Rules:    Used:      0 Available:   256
Virtual Slice * (physical slice 3) Rules:    Used:      0 Available:   256
Stage: LOOKUP
Slices:
Used: 0 Available: 4
Virtual Slice * (physical slice 0) Rules:    Used:      0 Available:   256
Virtual Slice * (physical slice 1) Rules:    Used:      0 Available:   256
Virtual Slice * (physical slice 2) Rules:    Used:      0 Available:   256
Virtual Slice * (physical slice 3) Rules:    Used:      0 Available:   256
```

```
Stage: EXTERNAL

Virtual Slice :  (*) Physical slice not allocated to any virtual slice.
```

History

This command was first available in ExtremeXOS 11.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show access-list width

```
show access-list width [slot slotNo | all]
```

Description

Displays the wide ACL mode configured on the supported switch or slot.

Syntax Description

<i>slotNo</i>	Specifies the slot to display.
all	Specifies all slots.

Default

N/A.

Usage Guidelines

Use this feature to display the width of the ACL TCAM key configured on a switch as being double wide or single wide.

Example

The following command displays the wide key mode on all slots:

```
show access-list width slot all
```

Following is sample output for this command:

Slot	Type	Width (Configured)
----	-----	-----
1	X460G2-48t-10G4	Single
2	X460G2-48t-10G4	Single
3	X670G2-48x-4q	Double
4		Single
5		Single
6		Single

7	Single	
8	Single	double

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

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.

This 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**—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.
- **Accounts locked out**—Account configured to be locked out after three consecutive failed login attempts (using the `configure account password-policy lockout-on-login-failures` command).



Note

This command does not show the failsafe account.

Example

The following command displays user account information on the switch:

```
show accounts pppuser
```

Output from this command looks similar to the following:

User Name	Access	LoginOK	Failed
admin	R/W	3	1
user	RO	0	0
dbackman	R/W	0	0
ron*	RO	0	0
nocteam	RO	0	0

(*) - Account locked

The following command displays the lawful intercept account distinguished by the "R/L" displayed in the Access column:

```
* (Private) X440e-24t.9 # show accounts
      User Name      Access LoginOK  Failed
-----
      admin          R/W        6        0
      user            RO         0        0
      myLIuser        R/L        N/A       N/A
```

History

This command was first available in ExtremeXOS 11.0.

Lawful intercept output was added in ExtremeXOS 15.3.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show accounts password-policy

```
show accounts password-policy
```

Description

Displays password policy information for all users on the switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

To view the password management information, you must have administrator privileges.

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

- Global password management parameters applied to new accounts upon creation:
 - Maximum age—The maximum number of days for the passwords to remain valid.
 - History limit—The number of previous password that the switch scans prior to validating a new password.
 - Minimum length—The minimum number of characters in passwords.
 - Character validation—The passwords must be in the specific format required by the `configure account password-policy char-validation` command.
 - Lockout on login failures—If enabled, the system locks out users after 3 failed login attempts.
 - Accounts locked out—Number of accounts locked out.
- **User Name**—The name of the user. This list displays all of the users who have access to the switch.
- **Password Expiry Date**—Date the password for this account expires; may be blank.
- **Password Max. age**—The number of days originally allowed to passwords on this account; may show None.
- **Password Min. length**—The minimum number of characters required for passwords on this account; may show None.
- **Password History Limit**—The number of previous passwords the system scans to disallow duplication on this account; may show None.

Example

The following command displays the password management parameters configured for each account on the switch:

```
# show accounts password-policy

-----
Accounts global configuration(applied to new accounts on creation)
-----
Password Max. age           : None
Password Min. age           : None
Password Min. Different Chars : 8
Password History limit      : None
Password Min. length        : None
Password Character Validation : Disabled
Accts. lockout on login failures: Disabled
Lockout time period         : Until Cleared
-----

-----
User Name      Expiry      Max. Expiry      Min. Min. Min. Hist
Lockout Lockout Flags      Date      Age  Date      age  len  diff Limit
Time      Time
              (Max)              (Min)              char
Config  Remain
-----
-----
None      U      - ---      admin      None      None None 0
```



```

None      U      - ---      user      None      None None      0
-----
Lockout Time Config: (U) Account is locked until cleared via 'clear account <name>
lockout'.
Flags: (C) Password character validation enabled, (L) Account locked out,
      (l) Account lockout on login failures enabled

```

History

This command was first available in ExtremeXOS 11.2.

Minimum different characters for changed password and minimum lifespan for passwords information was added in ExtremeXOS 30.7.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show auto-peering

```
show auto-peering {bgp | ospf}
```

Description

This command displays the status of BGP or OSPF auto-peering and the learned auto-peering interfaces and corresponding remote peer information.

Syntax Description

ospf	Shows OSPF auto-peering status and interface information.
bgp	Shows BGP auto-peering status and interface information.

Default

N/A

Usage Guidelines

All existing BGP show commands can be used to display the status of BGP peers and routes.

Example

The following example shows the status of BGP auto-peering:

```

# show auto-peering bgp
Type      : BGP
Password  :
Id        : 0
Router ID: 2.2.2.102

```

```

AS      : 102

Peer Id      Password
-----
2222      None
2223      #75Zvb3YfCBE/4+eSQE5dA5T2lmhF5A==

VLAN      Neighbor IP Address      Router ID      Remote AS      Port      Fabric ID
Peer
-----
SYS_BGP_0002  fe80::204:96ff:fe9d:66e8  3.3.3.103      103      69      0      Yes

```

The following example shows OSPF auto-peering information:

```

# show auto-peering ospf

Type      : OSPF
Id        : 0
Router ID: 1.1.1.54

VLAN      Router ID      Port      Fabric ID      Peer
-----
SYS_OSPF_0002  1.1.1.55      19      0      No

```

History

This command was first available in ExtremeXOS 22.5.

Peer ID and password information was added in ExtremeXOS 30.5.

The **ospf** option was added in ExtremeXOS 30.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

This feature requires the Advanced Edge license. For more information about licenses, see the [Switch Engine 32.3 Feature License Requirements](#).

show auto-peering one-config

```

show auto-peering one-config {remote | service | route | database |
bootprelay}

```

Description

Shows auto-peered OneConfig information configured on the device.

Syntax Description

auto-peering	Specifies Auto-peering.
one-config	Specifies changing the BGP Auto-peering OneConfig configuration.
remote	Displays only remote ID information.

service	Displays only service information.
route	Displays only host-driven static route information.
database	Displays only overlay database information.
bootprelay	Displays only BOOTP relay information.

Default

N/A.

Usage Guidelines

To configure remote IDs, use the configure **auto-peering one-config remote id** [**add** | **delete**] *id* {**password** [**none** | {**encrypted**} *tcpPassword*]} command.

To configure services, use the configure **auto-peering one-config nsi-id** *id* **type** [**nsi** | **vrf**] [**add** | **delete**] [[*ipaddress* {*netmask*} | *ipNetmask*] | *ipNetmaskv6*] {**vr** *vrname* } command.

To configure routes, use the configure **auto-peering one-config iproute** [**add** | **delete**] [*host* | *hostv6*] [[*ipaddress* {*netmask*} | *ipNetmask*] *gateway* | *ipNetmaskv6* *gatewayv6*] {**vr** *vrname* } command.

To configure overlay databases, use the configure **auto-peering one-config overlay** [**add** | **delete**] **server** [*address* | *addressv6*] {**type** **bgp-rr** } {**id** *id*} {**password** [**none** | {**encrypted**} *password*]} command.

To configure BOOTP relay, use the configure **auto-peering one-config bootprelay** [**add** | **delete**] [*ip_address* | *ipv6_address*] **vr** *vrname* comamnd.

Example

The following example shows all auto-peered OneConfig information:

```
# show auto-peering one-config
ID                : 0
Password          : None
Route Target      : None
Anycast MAC       : 00:00:00:1b:1b:1b

OneConfig Allowed Remote IDs

Remote ID  Password
-----
23456      None
34567      #HktHhKOfo/Tk1YtJC1pBz24ZMBHCGg==

OneConfig Dynamic Service

NSI Service  Type    VRF      Address                Installed
-----
1000         VRF     red      0.0.0.0/0              YES
```

```

1001      NSI      red      50.1.101.1/24      YES
1002      NSI      red      50.1.102.1/24      YES
1003      NSI      red      50.1.103.1/24      NO
1004      NSI      red      50.1.104.1/24      YES
2000      VRF      blue      0.0.0.0/0          YES
2001      NSI      blue      50.1.201.1/24      YES
2002      NSI      blue      50.1.202.1/24      YES
2003      NSI      blue      50.1.203.1/24      NO
2004      NSI      blue      50.1.204.1/24      YES
3001      NSI      VR-Default 50.1.131.1/24      YES
3002      NSI      VR-Default 50.1.132.1/24      YES
3003      NSI      VR-Default 50.1.133.1/24      NO
4002      NSI      VR-Default 1000:2000:3000:4000:5000:6000:7000:8000/128 NO

```

OneConfig Dynamic Host Static Routes

VRF	Host	Route	Gateway
Installed			
-----	-----	-----	-----
red	50.1.102.101	60.1.1.0/24	50.1.102.101
NO			
red	50.1.103.104	70.1.1.0/24	50.1.103.104
NO			
red	50.1.103.104	80.1.1.0/24	50.1.103.104
NO			

OneConfig Overlay Database

Database Address	Type	ID	Password
-----	-----	-----	-----
50.1.133.104	Redis	2000	None
50.1.133.105	BGP_RR	2000	None

OneConfig Dynamic Bootprelay

VRF	Server
-----	-----
VR-Default	50.1.131.105
blue	50.1.201.105
red	50.1.101.105

History

This command was first available in ExtremeXOS 31.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show auto-provision

```
show auto-provision {{vr} vr_name}
```

Description

Displays the current state of auto provision on the switch.

Syntax Description

<i>vr_name</i>	Specifies the virtual router. This may be <i>VR-Default</i> or <i>VR-Mgmt</i> only
----------------	--

Default

N/A.

Usage Guidelines

Use this command to display the current state and the statistics of the auto provision feature on the switch.

Example

The following command displays all information on the current state of auto provision:

```
show auto-provision
```

Following is sample output for the command when the auto provision is enabled. When “Enabled” the feature can be “In progress”, “Done”, or “Failed.”

```
(Auto-Provision) switch # show auto-provision
-----
VR-Name      Auto-Provision Status  Number of attempts
-----
VR-Default   Enabled (In progress)  2
VR-Mgmt      Enabled (In progress)  1
switch # show auto-provision
-----
VR-Name      Auto-Provision Status  Number of attempts
-----
VR-Default   Enabled (Done)         0
VR-Mgmt      Enabled (Done)         0
```

The following command displays information on the current state of auto provision on VR-Mgmt.

```
show auto-provision vr "VR-Mgmt"
```

Following is sample output for the command when auto provision is disabled:

```
switch # show auto-provision vr "VR-Mgmt"
DHCP Auto-Provision      : Disabled
Number of attempts       : 0
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show automation edge database

`show automation edge database`

Description

Displays Automation Edge remote VXLAN network identifier (VNI)-device database information.

Syntax Description

automation	Displays Automation Edge VXLAN VNI-device database information.
edge	Specifies Automation Edge.
database	Specifies Automation Edge database information.

Default

N/A.

Usage Guidelines

N/A.

Example

The following example shows database information:

```
# show automation edge database
Database Name  Status
-----
SampleDB      Up

Database Name  Status
-----
SampleDB      Down
```

History

This command was first available in ExtremeXOS 31.1 as a demonstration feature.

This command is fully supported in ExtremeXOS 31.2.

Platform Availability

This command is available on the ExtremeSwitching 5520 platform.

show automation edge devices

```
show automation edge devices {vni vni}
```

Description

Displays network devices information for an Automation Edge remote VXLAN network identifier (VNI)-device database.

Syntax Description

automation	Displays Automation Edge VXLAN VNI-device database information.
edge	Specifies Automation Edge.
devices	Specifies network devices information, such as switches or APs.
vni	A unique 24-bit Virtual Network Identifier";%
<i>vni</i>	Virtual Network Identifier value between 1 and 16777215";

Default

N/A.

Usage Guidelines

This command shows the details about the network devices, including device identifier, IP address of the device, MAC address of the device, and the device type (access point or switch).

Example

The following example shows network devices information:

# show automation edge devices				
VNI	Device Name	Serial Number	IP Address	Type

30	APName1	1111222233334444_110	200.200.200.10	AP
31	APName2	1111222233334444_11	200.200.200.1	AP

The following example shows network devices information for VNI "100":

# show automation edge devices vni 100			
Device Name	Serial Number	IP Address	Type

APName1	1111222233334444_15	200.200.200.5	AP
APName2	1111222233334444_17	200.200.200.7	

History

This command was first available in ExtremeXOS 31.1 as a demonstration feature.

This command is fully supported in ExtremeXOS 31.2.

Platform Availability

This command is available on the ExtremeSwitching 5520 platform.

show avb

show avb

Description

Displays a summary of MSRP, MVRP, and gPTP configuration on the switch.

Syntax Description

avb	Audio Video Bridging.
------------	-----------------------

Default

N/A.

Usage Guidelines

Use this command to display a summary of MSRP, MVRP, and gPTP configuration and status on the switch.

Example

```
#show avb
gPTP status           : Enabled
gPTP enabled ports    : *17d *19d

MSRP status           : Enabled
MSRP enabled ports    : !3 *17ab *19a

MVRP status           : Enabled
MVRP enabled ports    : *17 *19

Flags:                (*) Active, (!) Administratively disabled,
                      (a) SR Class A allowed, (b) SR Class B allowed,
                      (d) Disabled gPTP port role, (m) Master gPTP port role,
                      (p) Passive gPTP port role, (s) Slave gPTP port role.
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show bandwidth pool

show bandwidth pool [**ingress** | **egress** | **duplex**] **vlan** *vlan_name*

Description

Displays the configured bandwidth pool settings for the specified VLAN.

Syntax Description

ingress	Displays configured bandwidth pool settings for incoming traffic only.
egress	Displays configured bandwidth pool settings for outgoing traffic only.
duplex	Displays configured bandwidth pool settings for traffic in both directions.
<i>vlan_name</i>	Displays configured bandwidth pool settings only for the specified VLAN.

Default

N/A.

Usage Guidelines

This command displays the configured bandwidth pool settings for a VLAN. Values displayed include the VLAN, maximum reserveable bandwidth (both ingress and egress), and bandwidth reserved by application and by priority level.

Example

The following command displays bandwidth pool settings and accepted bandwidth reservations for all ports:

```
show bandwidth pool duplex vlan vlan_1
# show bandwidth pool duplex vlan vlan_1
(mbps)  Rsvd CIRBW  Cmnt  Cmnt CIRBW
Vlan      Dir   Phy   BE Limit Pools Total Avail
-----
vlan_1      Rx  1000      0  1000   300   300   700
Tx  1000      0  1000   500   500   500
-----
(mbps)  CIRBW Available in Pool (per priority level)
Appl  Dir  Pool      0      1      2      3      4      5      6      7
-----
mpls   Rx   300   300   300   290   290   290   290   290   290
Tx   500   500   500  491   491   491   491   491   491
(Rx)-Receive, (Tx)-Transmit  (BE)-Best Effort
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

show banner

```
show banner { after-login | before-login }
```

Description

Displays the user-configured banners.

Syntax Description

after-login	Specifies the banner that is displayed after login.
before-login	Specifies the banner that is displayed before login.

Default

N/A.

Usage Guidelines

Use this command to display specific configured CLI banners.

If no keywords are specified, all configured banners are displayed. To display a specific banner, use the before-login or after-login keyword.

Example

The following command displays the configured CLI switch banners:

```
show banner
```

Output from this command varies depending on your configuration; the following is one example:

```
Before-login banner:
Extreme Networks Summit Switch
#####
Unauthorized Access is strictly prohibited.
Violators will be prosecuted
#####
Acknowledge: Enabled
After-login banner:
Press any key to continue
```

History

This command was first available in ExtremeXOS 10.1.

The after-login option was added in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show banner netlogin

```
show banner netlogin
```

Description

Displays the user-configured banner string for network login.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to view the banner that is displayed on the network login page.

Example

The following command displays the network login banner:

```
show banner netlogin
```

If a custom banner web page exists, show banner netlogin generates the following output:

```
***** Testing NETLOGIN BANNER at <system name>*****  
NOTE: Banner is not in use. Overridden since custom login page  
"netlogin_login_page.html" is present.
```

If a custom banner web page does not exist, show banner netlogin generates the following output:

```
***** Testing NETLOGIN BANNER at <system name>*****
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show bfd

```
show bfd
```

Description

Displays information on existing BFD sessions.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to show the status of the current BFD sessions.

The following session states are displayed:

- Init—The state when BFD is establishing the session.
- Down—The state when BFD detects that the session is down.
- Admin Down—The state when the user disables BFD on that interface.
- Up—The state when the BFD session is established.

Example

The following command displays information on current BFD sessions:

```
# show bfd
```

Following is sample output from this command:

```
Number of sessions           : 2
Sessions in Init State       : 0
Sessions in Down State       : 0
Sessions in Admin Down State : 1
Sessions in Up State         : 1

SNMP Traps for Session Down  : Enabled
SNMP Traps for Session Up    : Enabled
SNMP Traps Batch Delay       : 1000 ms
Hardware Assist Operational State : Enabled
                               (or)
                               :
Disabled                     (Incapable heterogeneous stack)
                               (Incapable standalone platform)
                               (Loopback port not configured)

Hardware Assist Primary Loopback Port : 1:1
Hardware Assist Secondary Loopback Port : None
Maximum # of Hardware Assist Sessions : 2047
```

History

This command was first available in ExtremeXOS 12.4.

The hardware assist output was added in ExtremeXOS 21.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show bfd counters

```
show bfd counters
```

Description

Displays the readings of the global BFD counters.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to display global BFD counters.

To clear the counters, use the `clear counters bfd` command.

Example

The following command displays BFD global counters:

```
# show bfd counters

Valid Tx Pkt           : 177      Valid Rx Pkt           : 177
Rx Invalid TTL         : 0        Rx Invalid UDP SrcPort  : 0
Interface Not found    : 0        Rx Invalid Version     : 0
Rx Invalid Length Pkt  : 0        Rx Invalid Multiplier   : 0
Rx Invalid Demand Mode : 0        Rx Poll & Final set     : 0
Rx Invalid My Discriminator : 0    Rx Invalid Your Discriminator : 0
Rx Invalid Auth Length : 0        Rx session Not Found    : 6
Auth Type Fails        : 0        Authentication Fails     : 0
```

```
Tx Fails           : 0      Rx Discarded Pkt      : 0
Rx Invalid Multipoint : 0
```



Note

The Rx session Not Found counter is incremented when the BFD session corresponding to the received BFD packet is not found. The Rx Discarded Pkt counter is incremented when the neighbor state indicated in the BFD packet is not one of the expected/allowed states.

History

This command was first available in ExtremeXOS 12.4.

Rx Invalid Multipoint information was added in ExtremeXOS 30.7.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show bfd session client

```
show bfd session client [ bgp {ipv4 | ipv6} | mpls | ospf {ipv4 | ipv6}
                        | static {ipv4 | ipv6}] {vr [vrname | all]}
```

Description

Displays the BFD session information for a specified client.

Syntax Description

bgp	<i>BGP</i>
mpls	Specifies an <i>MPLS</i> client.
ospf	<i>OSPF</i> protocol.
ipv4	Displays sessions requested by IPv4 version client, e.g. OSPFv2 (Default).
ipv6	Displays sessions requested by IPv6 version client, e.g. <i>OSPFv3</i> .
static	Specifies a static route.
<i>vrname</i>	Specifies the name of the virtual router.

Default

IPv4.

Usage Guidelines

Use this command to display session information for a specified client.

Example

The following command displays the BFD sessions for an MPLS client on all VRs:

```
# show bfd session client mpls vr all
```

Following is sample output from this command:

```
Neighbor      Interface      Detection      Status
-----
10.10.10.2    vlan10         3000           Up
=====
NOTE: All timers in milliseconds.
```

History

This command was first available in ExtremeXOS 12.4.

Support for BFD protected static route was added in ExtremeXOS 12.5.3.

The **ospf** keyword was added in ExtremeXOS 15.3.2.

Support for border gateway protocol (BGP) was added in ExtremeXOS 21.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show bfd session counters missed-hellos

```
show bfd session counters missed-hellos {session-id first {- last} |
    neighbor ipaddress {vr [vrname | all]} | vr [vrname| all]} {detail |
    no-refresh | refresh}
```

Description

This command displays statistics of missed hello packets.

Syntax Description

session-id	Display statistics for sessions having session ID within the given range.
<i>first</i>	Only or first of range of session
<i>last</i>	Last of range of session ID.
<i>ipaddress</i>	Specify IPv4 or IPv6 Destination address.
vr	Virtual router.
<i>vrname</i>	Virtual router name.
all	All virtual routers.
detail	Detailed view of statistics.

no-refresh	Page by page display without continuous refresh.
refresh	Continuous refresh of output.

Default

Refresh.

Usage Guidelines

You can select the sessions by either neighbor IP address, by range of session IDs, by VR or display all the available sessions. Display selection by session ID is useful if the neighbor IP is link-local and VLAN name is long (i.e. close to 32 characters).

Example

The following example displays summary view with the refresh option.

```
# show bfd session counters missed-hellos
=====
Neighbor                               Session ID  Number Of
Misses                                1      2      2+
=====
fe80::204:96ff:fe7e:c2f%test           251    15     8     7
fe80::204:96ff:fe7e:c2f%verify         252    10     6     4
50.0.0.1                               300 >9999 >9999 >9999
```

History

This command was first available in ExtremeXOS 21.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show bfd session counters vr all

```
show bfd session {ipv4 | ipv6} {ipaddress} counters {vr [vrname | all]}
```

Syntax Description

ipv4	Displays all IPv4 sessions.
ipv6	Displays all IPv6 sessions.
<i>ipaddress</i>	Displays sessions in specified VR.

Default

Displays all IPv4 sessions counters by default if IPv4 or IPv6 is not specified.

Usage Guidelines

Use this command to display BFD session counters.

To clear the counters, use the `clear counters bfd` command.

Example

The following command displays the session counters:

```
# show bfd session counters vr all
```

Following is sample output from this command:

```
Neighbor : 10.10.10.1      Interface : vlan10Vr-Name :   bfd_vr10
Valid Rx Pkt              : 87
Total Tx Pkt              : 87
Auth Type Fails           : 0
Authentication Fails      : 0
Discarded Pkt             : 0
```

History

This command was first available in ExtremeXOS 12.4.

IPv6 version of this command was added in ExtremeXOS 15.3.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show bfd session detail vr all

```
show bfd session {ipv4 | ipv6} {ipaddress } detail {vr [vrname | all]}
```

Description

Displays detailed information about a BFD session.

Syntax Description

ipv4	Displays all IPv4 sessions.
ipv6	Displays all IPv6 sessions.
<i>ipaddress</i>	Displays sessions in specified VR.
vrname	Displays sessions in specified VR.

Default

Displays all IPv4 sessions by default if `ipv4` or `ipv6` is not specified.

Usage Guidelines

Use this command to display BFD session information in detail.

Example

The following command displays the BFD session information in detail:

```
show bfd session detail vr all
```

Following is sample output from this command:

```
# show bfd session detail vr all
Neighbor      : 10.10.10.1          Local      : 10.10.10.2
Vr-Name       : bfd_vr10           Interface   : vlan10
Session Type  : Single Hop         State       : Up
Detect Time   : 3000 mc            Age         : 250 ms
Discriminator (local/remote)      : 1 / 1
Demand Mode (local/remote)        : 0 / 0
Poll (local/remote)               : 0 / 0
Tx Interval (local/remote)        : 1000 / 1000 ms
Rx Interval (local/remote)        : 1000 / 1000 ms
oper Tx Interval                   : 1000 ms
oper Rx Interval                   : 1000 ms
Multiplier (local/remote)         : 3 / 3
Local Diag                        : 0 (No Diagnostic)
Remote Diag                       : 0 (No Diagnostic)
Authentication                     : None
Clients                           : MPLS,
Uptime                            : 00 days 00 hours 00 minutes 41 seconds
Up Count                          : 1
Last Valid Packet Rx              : 00:51:49.300000
Last Packet Tx                    : 00:51:48.820000
```

The following command displays a specified IPv6 BFD session in detail:

```
# show bfd session fe80::204:96ff:felf:a800%v2 detail

Neighbor      : fe80::204:96ff:felf:a800
Local         : fe80::204:96ff:fe27:2c6a
VR-Name       : VR-Default          Interface   : v2
Session Type  : Single Hop         State       : Up
Detect Time   : 60000 ms           Age         : 460 ms
Discriminator (local/remote)      : 1 / 1
Demand Mode (local/remote)        : Off / Off
Poll (local/remote)               : Off / Off
Tx Interval (local/remote)        : 20000 / 1000 ms
Rx Interval (local/remote)        : 20000 / 1000 ms
Oper Tx Interval                   : 20000 ms
Oper Rx Interval                   : 20000 ms
Multiplier (local/remote)         : 3 / 3
Local Diag                        : 0 (No Diagnostic)
Remote Diag                       : 0 (No Diagnostic)
Authentication                     : None
Clients                           : OSPFv3
Uptime                            : 00 days 01 hours 35 minutes 43 seconds
Up Count                          : 9
Last Valid Packet Rx              : 12:27:36.464105
Last Packet Tx                    : 12:27:19.34236
```

History

This command was first available in ExtremeXOS 12.4.

IPv6 version was added in ExtremeXOS 15.3.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show bfd session vr all

```
show bfd session {ipv4 | ipv6} {ipaddress } { vr [vrname | all ] }
```

Description

Displays general information about a BFD session.

Syntax Description

ipv4	Displays all IPv4 sessions.
ipv6	Displays all IPv6 sessions.
<i>ipaddress</i>	Displays session that has specified address as destination address.
vrname	Displays sessions in specified VR.

Default

Displays all IPv4 sessions by default if ipv4 or ipv6 keyword is not specified.

Usage Guidelines

Use this command to display general information about a BFD session.

Example

The following command displays general information about the BFD session:

```
# show bfd session vr all
```

Following is sample output from this command:

```
Neighbor      Interface      Clients  Detection  Status      VR
=====
30.30.30.2    bfdVlan        ----s    0          Down        VR-Default
=====
Clients Flag: b - BGP, m - MPLS, o - OSPF, s - Static
NOTE: All timers in milliseconds.
```

Following is sample output with hardware assist information displayed:

```
)
# show bfd session detail vr all
Neighbour      : 10.10.10.1          Local      : 10.10.10.2
Vr-Name        : bfd_vr10         Interface  : bfd_vlan10
Session Type   : Single Hop    State      : Up
```

```

...
Up Count                : 1
Last Valid Packet Rx    : 00:51:49.300000
Last Packet Tx          : 00:51:48.8200000
Hardware Assist         : Yes

Neighbour      : 10.10.11.1      Local      : 10.10.11.2
Vr-Name       : bfd_vr10        Interface  : bfd_vlan11
Session Type  : Single Hop      State      : Up
...
Up Count                : 1
Last Valid Packet Rx    : 00:51:49.300000
Last Packet Tx          : 00:51:48.8200000
Hardware Assist         : Yes

```

History

This command was first available in ExtremeXOS 12.4.

The hardware assist example output was added in ExtremeXOS 21.1.

Support for *BGP* was added in ExtremeXOS 21.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show bfd vlan counters

```
show bfd vlan {vlan_name} counters
```

Description

Displays BFD counters on a specified *VLAN*.

Syntax Description

<i>vlan_name</i>	Specifies the VLAN name.
------------------	--------------------------

Default

N/A.

Usage Guidelines

Use this command to display counter readings for a specified VLAN.

This command is not supported by hardware-assisted BFD.

Example

The following command displays the counter readings for the VLAN `vlan10`:

```
# show bfd vlan vlan10 counters
```

Following is sample output from this command:

```
VLAN                               : vlan10
Valid Rx Pkt                       : 144
Total Tx Pkt                       : 144
Auth Type Fails                    : 0
Authentication Fails               : 0
Discarded Pkt                      : 0
Rx session Not Found               : 6
```



Note

The Discarded Pkt counter is incremented when the neighbor state indicated in the BFD packet is not one of the expected/allowed states. The Rx session Not Found counter is incremented when the BFD session corresponding to the received BFD packet is not found.

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show bfd vlan

```
show bfd vlan {vlan_name}
```

Description

Displays the BFD settings for the specified VLAN.

Syntax Description

<code> vlan_name </code>	Specifies the VLAN name.
--------------------------	--------------------------

Default

N/A.

Usage Guidelines

Use this command to display the BFD settings on a specified VLAN.

Example

The following command displays the BFD settings for the VLAN vlan10:

```
# show bfd vlan vlan10
```

Following is sample output from this command:

```
VLAN           : vlan10
BFD            : Enabled
Tx Interval    : 1000
Rx Interval    : 1000
Detection Multiplier : 3
Authentication : None
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show bgp

```
show bgp
```

Description

Displays *BGP* configuration information.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command examples display various BGP configurations:

Output for `show bgp` for a VRF (PE-CE Protocol, RD and RT configured):

```
(virtual-router vrf-foo) # show bgp
Enabled           : No           OperStatus         :           Down
```

```

RouterId      : 3.3.3.3      AS      : 200
LocalPref     : 100          MED      : None
Always-Compare-MED : Disabled Aggregation : Disabled
Route Reflector : No         RR ClusterId : 0
IGP Synchronization : Disabled New Community Format: Disabled
Routes from EBGp : 2         Routes from IBGP : 0
Routes redistributed: 1       Out Updates queued : 0
Fast Ext Fallover : Disabled MPLS LSP as Next-Hop: No
AS Disp Format  : Asplain     Maximum ECMP Paths : 1
ConfedId       : 0           Multipath-Relax : Disabled
Confed Peers   :
Networks       : 2
  ipv4-unicast 10.0.0.0/16 network-policy nwk.pol
  ipv4-multicast 11.0.0.0/16 network-policy nwk.pol
Aggregate Networks : 2
  ipv4-unicast 21.0.0.0/8 as-match advertise-policy: agg.pol
  ipv4-multicast 22.0.0.0/8 as-set summary-only advertise-policy: agg.pol
Route Statistics:
Address family      EBGp    IBGP    Redist.
-----
ipv4-unicast        0       0       0
ipv4-multicast      0       0       0

Redistribute       :
-----
Address Family

  Route Type      Flags      Priority      Policy
-----
ipv4-unicast

  Direct          EO          2048         None

ipv6-multicast

  Direct          EO          2048         None
-----
Flags: (E) Export Enabled, (L) Export Operationally Off due to Low Memory,
      (O) Export Operationally On

Advertise Inactive Routes:
ipv4-unicast : Disabled
ipv4-multicast : Disabled
ipv6-unicast : Disabled
ipv6-multicast : Disabled
ipv4-vxlan : Disabled

```

Output of show bgp for a VRF (PE-CE Protocol, RD and RT “not” configured):

```

# show bgp
Enabled      : No          OperStatus   : Down
RouterId     : 3.3.3.3     AS           : 200
LocalPref    : 100        MED          : None
Always-Compare-MED : Disabled Aggregation : Disabled
Route Reflector : No      RR ClusterId : 0
IGP Synchronization : Disabled New Community Format: Disabled
Fast Ext Fallover : Disabled MPLS LSP as Next-Hop: No
AS Disp Format  : Asplain   Maximum ECMP Paths : 1
ConfedId       : 0         Outbound rt. filter : Enabled
Confed Peers   :
Networks       : 4
  ipv4-unicast 10.0.0.0/16 network-policy nwk.pol

```

```

ipv4-multicast 11.0.0.0/16 network-policy nwk.pol
ipv6-unicast 2001::/64 network-pol nwk6.pol
ipv6-multicast 2001::/64 network-pol nwk6.pol
Aggregate Networks : 4
  ipv4-unicast 21.0.0.0/8 as-match advertise-policy: agg.pol
  ipv4-multicast 22.0.0.0/8 as-set summary-only advertise-policy: agg.pol
  ipv6-unicast 2003::/64 as-match advertise-policy: agg6.pol
  ipv6-multicast 2004::/64 as-set advertise-policy: agg6.pol

Route Statistics:
  Address family      EGBP    IBGP    Redist.
  -----
  ipv4-unicast        0        0        0
  ipv4-multicast      0        0        0
  ipv6-unicast        0        0        0
  ipv6-multicast      0        0        0

Redistribute:

-----
Address Family
  Route Type      Flags      Priority      Policy
  -----
  ipv4-unicast
    Direct        EO          2048          None
  ipv6-multicast
    Direct        EO          2048          None
  -----

Flags: (E) Export Enabled, (L) Export Operationally Off due to Low Memory,
       (O) Export Operationally On

Advertise Inactive Routes:
  ipv4-unicast : Disabled
  ipv4-multicast : Disabled
  ipv6-unicast : Disabled
  ipv6-multicast : Disabled
  ipv4-vxlan : Disabled

```

If BGP is added as a protocol inside a heavy-weight VR, normal BGP peering applies with the addition of vpnv4 address family support:

```

# show bgp
Enabled : No      OperStatus : Down
RouterId : 3.3.3.3 AS : 200
LocalPref : 100   MED : None
Always-Compare-MED : Disabled Aggregation : Disabled
Route Reflector : No RR ClusterId : 0
IGP Synchronization : Disabled New Community Format: Disabled
Fast Ext Fallover : Disabled MPLS LSP as Next-Hop: No
AS Disp Format : Asplain Maximum ECMP Paths : 1
ConfedId : 0 Outbound rt. filter : Enabled
Confed Peers :
Networks : 4
  ipv4-unicast 10.0.0.0/16 network-policy nwk.pol
  ipv4-multicast 11.0.0.0/16 network-policy nwk.pol
  ipv6-unicast 2001::/64 network-pol nwk6.pol
  ipv6-multicast 2001::/64 network-pol
    nwk6.pol
Aggregate Networks : 4
  ipv4-unicast 21.0.0.0/8 as-match advertise-policy: agg.pol
  ipv4-multicast 22.0.0.0/8 as-set summary-only advertise-policy: agg.pol
  ipv6-unicast 2003::/64 as-match advertise-policy: agg6.pol

```



```
ipv6-multicast 2004::/64 as-set advertise-policy: agg6.pol
```

Route Statistics:

Address family	EBGP	IBGP	Redist.

ipv4-unicast	0	0	0
ipv4-multicast	0	0	0
vpn4	0	0	0
ipv6-unicast	0	0	0
ipv6-multicast	0	0	0

Redistribute:

ipv4	Admin	Operational	Shutdown	Policy
unicast	Status	Status	Priority	

Direct	Disabled	Down	2048	None
Static	Disabled	Down	2048	None
RIP	Disabled	Down	2048	None
BlackHole	Disabled	Down	2048	None
OSPFIntra	Disabled	Down	2048	None
OSPFInter	Disabled	Down	2048	None
OSPFExt1	Disabled	Down	2048	None
OSPFExt2	Disabled	Down	2048	None
ISISL1	Disabled	Down	2048	None
ISISL2	Disabled	Down	2048	None
ISISL1Ext	Disabled	Down	2048	None
ISISL2Ext	Disabled	Down	2048	None

ipv4	Admin	Operational	Shutdown	Policy
multicast	Status	Status	Priority	

Direct	Disabled	Down	2048	None
Static	Disabled	Down	2048	None
RIP	Disabled	Down	2048	None
BlackHole	Disabled	Down	2048	None
OSPFIntra	Disabled	Down	2048	None
OSPFInter	Disabled	Down	2048	None
OSPFExt1	Disabled	Down	2048	None
OSPFExt2	Disabled	Down	2048	None
ISISL1	Disabled	Down	2048	None
ISISL2	Disabled	Down	2048	None
ISISL1Ext	Disabled	Down	2048	None
ISISL2Ext	Disabled	Down	2048	None

ipv6	Admin	Operational	Shutdown	Policy
unicast	Status	Status	Priority	

Direct	Disabled	Down	2048	None
Static	Disabled	Down	2048	None
Ripng	Disabled	Down	2048	None
Ospf3-intra	Disabled	Down	2048	None
Ospf3-inter	Disabled	Down	2048	None
Ospf3-extern1	Disabled	Down	2048	None
Ospf3-extern2	Disabled	Down	2048	None
ISISL1	Disabled	Down	2048	None
ISISL2	Disabled	Down	2048	None
ISISL1Ext	Disabled	Down	2048	None
ISISL2Ext	Disabled	Down	2048	None

ipv6	Admin	Operational	Shutdown	Policy
multicast	Status	Status	Priority	

Direct	Disabled	Down	2048	None
Static	Disabled	Down	2048	None

```

Ripng          Disabled Down      2048      None
Ospf3-intra    Disabled Down      2048      None
Ospf3-inter    Disabled Down      2048      None
Ospf3-extern1  Disabled Down      2048      None
Ospf3-extern2  Disabled Down      2048      None
ISISL1         Disabled Down      2048      None
ISISL2         Disabled Down      2048      None
ISISL1Ext      Disabled Down      2048      None
ISISL2Ext      Disabled Down      2048      None

Advertise Inactive Routes:
ipv4-unicast   : Disabled
ipv4-multicast : Disabled
ipv6-unicast   : Disabled
ipv6-multicast : Disabled
ipv4-vxlan     : Disabled

```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

This command was modified in ExtremeXOS 15.3 to reflect its operation in VRs and VRFs.

Support for IPv4 VXLAN was added in ExtremeXOS 22.3.

Support for BGP multipath-relax was added in ExtremeXOS 22.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

show bgp evpn

```
show bgp evpn
```

Description

Displays EVPN global configuration.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Example

The following example shows the EVPN global configuration for the switch:

```
# show bgp evpn
Enabled : Yes
OperStatus : Up
Local Identifier : 1.0.0.25 Restart Duration (sec) :
180
MAC Move Duration (sec) : 60 MAC Move
Limit : 10
MAC Withdraw Delay (sec) : 1
Ignore AS for Route Target Matching : On
```

History

This command was first available in ExtremeXOS 31.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature,see the [Switch Engine 32.3 Feature License Requirements](#) document.

show bgp evpn evi

```
show bgp evpn evi {evi-index evi_index} {vni vni}
```

Description

Show information about the EVPN instance table.

Syntax Description

bgp	Specifies BGP.
evpn	Specifies Ethernet VPN (RFC 7432).
evi	Shows the EVPN instance table.
evi-index	Specifies the EVI index.
vni	Specifies a particular virtual network identifier.
<i>vni</i>	Selects the VNI. Range is 1 to 16,777,215.

Default

N/A.

Usage Guidelines

Each EVPN instance represents a VLAN in a virtual network. In this implementation, there must be a 1-to-1 relationship between VLAN and VNI.

Example

The following example displays the currently active EVPN EVI instances:

```
# show bgp evpn evi
```

Flgs	EVI-Idx	Name	VNI	Local	VTEP	RD	Import	RTs	Export
*	190		10000	1.1.1.102	1.1.1.102:190		64716:16787216		
		64716:16787216							
*	1020		101020	1.1.1.102	1.1.1.102:1020		64716:16878236		
		64716:16878236							
*	3500		111103	1.1.1.102	1.1.1.102:3500		64716:16888319		
		64716:16888319							
*	4089		101021	1.1.1.102	1.1.1.102:4089		64716:16878237		
		64716:16878237							
*SM	125	MktEvi	777	1.1.1.102		65535:1	25:16777993		
		25:16777993							
								25:100	
		25:100							
SM		EngEvi	884	1.1.1.102		1:65535		20:1	
		20:1							

Flags: (*) Entry is active, (S) Entry is statically configured. (M) RD and RT are manually configured
RD: Route distinguisher
RT: Route target - See RFC 8365 section 5.1.2.1 for details on automatically generated RT decoding instructions.

Some static entries may have multiple route target sets. These are listed in rows immediately following the initial row for the EVI-Idx.

Total Displayed: 6 Static: 2 Dyn: 4

History

This command was first available in ExtremeXOS 30.2.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

show bgp evpn ipv4

```
show bgp evpn ipv4 {evi-index evi_index} {ip-address ip_address}
```

Description

Shows the IPv4 entries from the EVPN MAC/IP table.

Syntax Description

bgp	Specifies BGP.
evpn	Specifies Ethernet VPN (RFC 7432).
ipv4	Shows only the IPv4 entries from the EVPN MAC/IP table.
evi-index	Specifies the EVI index.
<i>evi_index</i>	Restricts the display to EVI index (should be equal to VLAN ID). Range is 1 to 4,094.
ip-address	Restricts the display to an IP address.
<i>ip_address</i>	Selects the IP address.

Default

N/A.

Usage Guidelines

This command allows you to view the current set of IPv4 addresses configured in EVPN. If the ESI and ESI-Port fields are non-zero, then the entry was learned over a shared interface. The remote LACP partner's 6-byte MAC address is part of the ESI. For a full decoding of the ESI, see *RFC 7432*. The source (Src) column indicates whether the entry was learned (L)ocally or (R)emotely. The local entries are from the IP ARP cache or a locally configured routable VLAN. The remote entries appear only if the "In Use" flag is set to yes.

Example

The following example shows the current set of IPv4 addresses configured in EVPN.

```
# show bgp evpn ipv4
```

Src	EVI-Idx	Destination	MAC	ESI-Port	BGP Next Hop	VNI
ESI					In Use	
L	190	192.168.190.102	00:04:96:9d:64:e2			10000
		00:01:02:03:04:05:06:07:08:09	22		Yes	
R	190	192.168.190.101	00:04:96:9d:66:e8		3.3.3.103	
	10000				Yes	
R	1020	20.1.1.3	00:04:96:9c:2c:a2		1.1.1.101	
	101020				Yes	
L	1020	20.1.1.1	00:04:96:9d:64:e2			
	101020				Yes	
R	1020	20.1.1.2	00:04:96:9d:66:e8		3.3.3.103	
	101020				Yes	
L	3500	1.1.103.1	00:04:96:9d:64:e2			
	111103				Yes	
R	4089	21.1.1.3	00:04:96:9c:2c:a2		1.1.1.101	
	101021				Yes	

```

L      4089 21.1.1.1      00:04:96:9d:64:e2
101021                               Yes
R      4089 21.1.1.2      00:04:96:9d:66:e8 3.3.3.103
101021                               Yes
R      4089 21.1.1.30     00:0f:20:98:87:5a 3.3.3.103
101021                               Yes

Src: (L) Local, (R) Remote
In Use: Yes/No - Indicates if entry is installed in IP ARP cache.

Total MAC/IP entries: 10

```

History

This command was first available in ExtremeXOS 30.2.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

show bgp evpn ipv6

```
show bgp evpn ipv6 {evi-index evi_index} {ip-address ip_address}
```

Description

Shows the IPv6 entries from the EVPN MAC/IP table.

Syntax Description

bgp	Specifies BGP.
evpn	Specifies Ethernet VPN (RFC 7432).
ipv6	Shows only the IPv6 entries from the EVPN MAC/IP table.
evi-index	Specifies the EVI index.
<i>evi_index</i>	Restricts the display to EVI index (should be equal to VLAN ID). Range is 1 to 4,094.
ip-address	Restricts the display to an IP address.
<i>ip_address</i>	Selects the IP address.

Default

N/A.

Usage Guidelines

This command allows you to view the current set of IPv6 addresses configured in EVPN. If the ESI and ESI-Port fields are non-zero, then the entry was learned over a shared interface. The remote LACP partner’s 6-byte MAC address is part of the ESI. For a full decoding of the ESI, see *RFC 7432*. The source (Src) column indicates whether the entry was learned (L)ocally or (R)emotely. The local entries are from the neighbor-discovery cache or a locally configured routable VLAN. The remote entries are seen only if the “In Use” flag is set to yes.

Example

The following example shows the current set of IPv6 addresses configured in EVPN.

```
# show bgp evpn ipv6

Src  EVI-Idx  Destination
      MAC                BGP Next Hop  VNI  ESI                                ESI-Port  In Use
R    1020  2022::2%tenant
      01:01:01:01:01:01  3.3.3.103          101020  00:01:02:03:04:05:06:07:08:09  22      Yes

Src: (L) Local, (R) Remote
In Use: Yes/No - Indicates if entry is installed in neighbor discovery cache.

Total MAC/IP entries: 1
```

History

This command was first available in ExtremeXOS 30.2.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

show bgp evpn l3vni

```
show bgp evpn l3vni {vr vr_name}}
```

Description

Shows the L3 VNI entries from the EVPN MAC/IP table.

Syntax Description

bgp	Specifies showing the BGP configuration.
evpn	Specifies showing the EVPN configuration.
l3vni	Specifies showing the integrated routing and bridging IP VRF VNI configuration.

vr	Specifies showing the VRF.
<i>vr_name</i>	Provides the VRF name.

Default

N/A.

Usage Guidelines

If you do not specify a VR, all VRs appear.

Example

The following example shows all L3 VNI entries:

```
1 # show bgp evpn l3vni
VRF Name                Layer 3 VNI  RD
-----
vr-a                    10300      2.2.2.2:10300
vr-b                    10100      2.2.2.2:10100
vr-c                    10200      2.2.2.2:10200
```

The following example shows L3 VNI entries for VR "vr-b":

```
# show bgp evpn l3vni vr vr-b
VRF Name                Layer 3 VNI  RD
-----
vr-b                    10100      2.2.2.2:10100
```

History

This command was first available in ExtremeXOS 30.6.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

show bgp evpn mac

```
show bgp evpn mac {mac-address mac_address}
```

Description

Shows the current set of MAC addresses configured in EVPN.

Syntax Description

bgp	Specifies BGP.
evpn	Specifies Ethernet VPN (RFC 7432).
mac	Shows only the MAC entries from the EVPN MAC/IP table.
mac-address	Specifies restricting the display to a particular MAC address.
mac_address	Selects the MAC address to show.

Default

N/A.

Usage Guidelines

If the ESI and ESI-Port fields are non-zero, then the entry was learned over a shared interface. The remote LACP partner's 6-byte MAC address is part of the ESI. For a full decoding of the ESI, see RFC 7432. The (S)ource column indicates whether the entry was learned (L)ocally or (R)emotely. The local entries are from the MAC forwarding database. The remote entries are in the MAC forwarding database only if the "In Use" flag is set to yes.

Example

The following example shows the current set of MAC addresses configured in EVPN.

```
# show bgp evpn mac
```

Src	EVI-Idx	MAC	BGP Next Hop	VNI	ESI
ESI-Port	In Use				
L	190	00:04:96:9d:64:e2		10000	00:01:02:03:04:05:06:07:08:09
22	Yes				
R	190	00:04:96:9d:66:e8	3.3.3.103	10000	Yes
R	1020	00:04:96:9c:2c:a2	1.1.1.101	101020	Yes
L	1020	00:04:96:9d:64:e2		101020	Yes
R	1020	00:04:96:9d:66:e8	3.3.3.103	101020	Yes
R	1020	01:01:01:01:01:01	3.3.3.103	101020	Yes
L	3500	00:04:96:9d:64:e2		111103	Yes
R	4089	00:04:96:9c:2c:a2	1.1.1.101	101021	Yes
L	4089	00:04:96:9d:64:e2		101021	Yes
R	4089	00:04:96:9d:66:e8	3.3.3.103	101021	Yes
R	4089	00:0f:20:98:87:5a	3.3.3.103	101021	Yes

Src: (L) Local, (R) Remote
In Use: Yes/No - Indicates if entry is installed in MAC forwarding database.

Total MAC/IP entries: 11

History

This command was first available in ExtremeXOS 30.2.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

show bgp memory

show bgp memory {**detail** | *memoryType*}

Description

Displays *BGP* specific memory usage.

Syntax Description

detail	Displays detail information.
<i>memoryType</i>	Specifies the memory type usage to display.

Default

N/A.

Usage Guidelines

To see the memory types that you can display, enter the show bgp memory command without any attributes.

Example

The following command displays detailed BGP output for a specific memory types:

```
Switch.16.3 # sh bgp memory
BGP Memory Information
-----
Current Memory Utilization Level:      GREEN
-----
Type                AN                AB
-----
Callbacks           1141             17039828
```

```

Buffers                19                8456
Memory Utilization Statistics:
-----
-----
Module Type           Module Id           MemType
Name                  Size      AN/AB/HN/HB
-----
PCT_NBASE_ROOT        0x000000000000    16777219
MEM_PROCESS_ENTRY      212      8/1696/8/1696
PCT_NBASE_ROOT        0x000000000000    16777230
MEM_NBB_DIAGS_BLOCK    3212     8/25696/8/25696
PCT_NBASE_ROOT        0x000000000000    16777233
MEM_UNFORMATTED        1508     1/1508/1/1508
PCT_NBASE_ROOT        0x000000000000    50528257    0x0003030001
                        732      1/732/1/732
PCT_NBASE_ROOT        0x000000000000    50921473
0x0003090001          2004     1/2004/1/2004
PCT_NBASE_ROOT        0x000000000000    52232193
0x00031d0001          1508     1/1508/1/1508
PCT_NBASE_ROOT        0x000000000000    1090584577
MEM_QBRM_LOCAL         9660     2/19320/2/19320
PCT_NBASE_ROOT        0x000000000000    1090650113
MEM_QBNM_LOCAL         1508     2/3016/2/3016
PCT_NBASE_ROOT        0x000000000000    1107361793
MEM_QVB_LOCAL          3076     1/3076/1/3076
PCT_SCK                0x0001109000    16777220
MEM_NBB_POOL_CB        108      9/972/9/972

PCT_QVB                0x0001104000    1107361803
MEM_QVB_RV_REM_CB      60       6/360/6/360
PCT_QVB                0x0001104000    1107361806
MEM_QVB_AS_PATH_CB     60       4/240/4/240
PCT_QVB                0x0001104000    1107361807
MEM_QVB_RTM_CB         516      1/516/1/516

Flags : AN - Number of Allocations,          AB - Total Allocation in
Bytes
      : HN - Number of High Water Marks for Allocation,  HB - Total High Water Mark
Allocations in Bytes
t16.3 # sh bgp memory 1107361807    BGP Memory Information
-----
-----
Current Memory Utilization Level:    GREEN
-----
-----
Type                AN                AB
-----
Callbacks           1141            17039828
Buffers              19              8456

Memory Statistics for MEM_QVB_RTM_CB:
-----
-----
MemId                Size                AN                AB
-----
001107361807         516                1                516

Flags : AN - Number of Allocations,          AB - Total Allocation in Bytes
      : HN - Number of High Water Marks for Allocation, HB - Total High Water Mark
Allocations in Bytes

```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

This command is updated to reflect L3 VPN changes in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

show bgp neighbor [flap-statistics | suppressed-routes]

For IPv4 and IPv6 address families:

```
show bgp {neighbor} remoteaddr {address-family [ipv4-unicast |
ipv4-multicast | ipv6-unicast | ipv6-multicast]} [flap-statistics
| suppressed-routes] {detail} [all | as-path path-expression |
community [no-advertise | no-export | no-export-subconfed | number
community-number | autonomous-system-id : bgp-community] | network
[any/netMaskLen | networkPrefixFilter] {exact}]
```

For the VPNv4 address family:

```
show bgp {neighbor} remoteaddr address-family vpnv4 [flap-statistics
| suppressed-routes] {detail} [all | as-path path-expression |
community [no-advertise | no-export | no-export-subconfed | number
community-number | autonomous-system-id : bgp-community] | rd
rd_value network [any/ netMaskLen | networkPrefixFilter] {exact}]
```

Description

Displays flap statistics or suppressed-route information about a specified neighbor.

Syntax Description

<i>remoteaddr</i>	Specifies an IPv4 or IPv6 address that identifies a BGP neighbor.
ipv4-unicast	Specifies the IPv4 unicast address family.
ipv4-multicast	Specifies an IPv4 multicast address family.
ipv6-unicast	Specifies the IPv6 unicast address family.
ipv6-multicast	Specifies an IPv6 multicast address family.
vpnv4	Specifies the VPNv4 address family for Layer 3 VPN support.
flap-statistics	Specifies that only flap-statistics should be displayed (for route flap dampening enabled routes).
suppressed-routes	Specifies that only suppressed routes should be displayed (for route flap dampening enabled routes).

detail	Specifies to display the information in detailed format.
all	Specifies all routes.
<i>path-expression</i>	Display routes that match the specified AA path expression.
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.
<i>community_number</i>	Specifies a community number.
<i>autonomous-system-id</i>	Specifies an autonomous system ID (0-65535).
<i>bgp-community</i>	Specifies the BGP community number.
rd	Specifies the Route Distinguisher (RD) value for the Layer 3 VPN routes for which you want to clear flap statistics.
any	Specifies all routes with a given or larger mask length.
<i>netMaskLen</i>	Specifies a IPv4 or IPv6 subnet mask length (number of bits).
<i>networkPrefixFilter</i>	Specifies an IPv4 or IPv6 address and netmask.
exact	Specifies an exact match with the IP address and subnet mask.

Default

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

This command applies to the current VR or VRF context.



Note

If this command displays Bad Source Address, the BGP neighbor IP address is unavailable. Possible causes for this condition include a deleted or unconfigured VLAN or IP address.

The option network any / *netMaskLen* displays all BGP routes whose mask length is equal to or greater than *maskLength*, irrespective of their network address.

The option network any / *netMaskLen* exact displays all BGP routes whose mask length is exactly equal to *maskLength*, irrespective of their network address.

If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.



Note

For an IPv6 peer, an IPv6 address family must be specified, because an IPv6 peer does not support IPv4 address families. If no address family is specified for an IPv6 peer, the default, the IPv4 unicast address family, applies and no address-family information appears. Similarly an IPv4 peer only supports IPv4 address families and no address-family information appears if an IPv6 address family is specified.

To display Layer 3 VPN information, you must enter this command in the context of on the *MPLS*-enabled VR; it is not supported for BGP neighbors on the CE (VRF) side of the PE router.

Example

The following command displays flap statistics for the specified IPv4 neighbor:

```
* Switch.18 # show bgp neighbor 10.0.0.0 flap-statistics
BGP Routes Flap Statistics
Destination          NextHop          Penalty Flaps Duration Reuse          AS-
Path
-----
* ?100:1:100.0.0.0/8  11.0.0.2          100
100
Flags: (*) Preferred BGP route, (>) Active, (d) Suppressed, (h) History
(s) Stale, (m) Multipath, (u) Unfeasible
Origin: (?) Incomplete, (e) EGP, (i) IGP
Total Number of Flapped Routes: 1
```

The following command displays flap statistics for the specified IPv6 neighbor:

```
* Switch.21 # show bgp neighbor 2001::64:: address-family ipv6-unicast flap-statistics
BGP Routes Flap Statistics
Destination          NextHop          Penalty Flaps Duration Reuse          AS-
Path
-----
* ?2001::/64          3001::1          100
100
Flags: (*) Preferred BGP route, (>) Active, (d) Suppressed, (h) History
(s) Stale, (m) Multipath, (u) Unfeasible
Origin: (?) Incomplete, (e) EGP, (i) IGP
Total Number of Flapped Routes: 1
```

History

This command was first available in ExtremeXOS 10.1.

The **any/netMaskLen** options were added in ExtremeXOS 11.0.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Support for Layer 3 VPNs was added in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

show bgp neighbor received orf

```
show bgp {neighbor} remoteaddr {address-family [ipv4-unicast | ipv4-multicast | vpnv4] } received-orf
```

Description

Displays on the remote speaker the ORF lists received and installed from the local speaker for installation and outbound route filtering for IPv4 and IPv6 address families.

Syntax Description

<i>remoteaddr</i>	Specifies an IPv4 address that identifies a <i>BGP</i> neighbor.
ipv4-unicast	Specifies IPv4 unicast routes.
ipv4-multicast	Specifies IPv4 multicast routes.
vpnv4	Specifies VPNv4 routes.
received-orf	Displays on the remote speaker the ORF lists it received, and subsequently installed from the local speaker for installation and outbound route filtering.

Default

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

ORF is only supported for IPv4 peers. If this command is executed for an IPv6 peer, the command is rejected with the following error message:

```
Outbound-route-filtering not supported for IPv6 peer remoteaddr
```

Example

The following example shows the ORF filters received by the remote speaker:

```
show bgp neighbor 11.0.0.2 received-orf
Address family: IPv4 unicast
Prefix list:
  nlri 21.0.0.0/8      exact      permit
  nlri 22.1.0.0/16    min 24 max 28    permit
  nlri 23.0.0.0/8      min 16          deny
```

```
Community list:

Extended-community list:
  rt:100:2          permit
  rt:100:3          permit
  rt:101:1          deny
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the L3 VPN feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

show bgp neighbor

For IPv4 and IPv6 address families:

```
show bgp {neighbor} remoteaddr {address-family [ipv4-unicast | ipv4-
multicast | ipv6-unicast | ipv6-multicast | ipv4-vxlan | {l2vpn-evpn
[inclusive-multicast | mac-ip | auto-discovery | esi | ip-prefix]]}]
[accepted-routes | received-routes | rejected-routes | transmitted-
routes] {detail} [all | as-path path-expression | community [no-
advertise | no-export | no-export-subconfed | number community_number
| autonomous-system-id : bgp-community] | network [any/netMaskLen |
networkPrefixFilter] {exact}]
```

For the VPNv4 address family:

```
show bgp {neighbor} remoteaddr address-family vpnv4 [accepted-routes
| received-routes | rejected-routes | transmitted-routes] {detail}
[all | as-path path-expression | community [no-advertise | no-
export | no-export-subconfed | number community_number | autonomous-
system-id :bgp-community] | rd rd_value network [any/netMaskLen |
networkPrefixFilter] {exact}]
```

Description

Displays information about routes to a specified neighbor.

Syntax Description

<i>remoteaddr</i>	Specifies an IPv4 or IPv6 address that identifies a <i>BGP</i> neighbor.
ipv4-unicast	Specifies IPv4 unicast routes.
ipv4-multicast	Specifies IPv4 multicast routes.
ipv6-unicast	Specifies IPv6 unicast routes.

ipv6-multicast	Specifies IPv6 multicast routes.
ipv4-vxlan	Specifies IPv4 VXLAN routes.
l2vpn-evpn	Specifies an L2VPN EVPN address family.
inclusive-multicast	Displays EVPN inclusive-multicast (type 3) routes
mac-ip	Displays EVPN MAC/IP (type 2) routes.
auto-discovery	Displays EVPN auto-discovery (type 1) routes.
esi	Displays EVPN Ethernet segment (type 4) routes.
ip-prefix	Displays EVPN ip-prefix (type 5) routes.
vpnv4	Specifies VPNv4 routes.
accepted-routes	Specifies that only accepted routes are displayed.
received-routes	Specifies that only received routes are displayed.
rejected-routes	Specifies that only rejected routes are displayed.
transmitted-routes	Specifies that only transmitted routes are displayed.
detail	Specifies to display the information in detailed format.
all	Specifies all routes.
<i>path-expression</i>	Display routes that match the specified AA path expression.
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.
<i>community_number</i>	Specifies a community number.
<i>autonomous-system-id</i>	Specifies an autonomous system ID (0-65535).
<i>bgp-community</i>	Specifies the BGP community number.
rd	Specifies the Route Distinguisher (RD) value for the Layer 3 VPN routes for which you want to clear flap statistics.
any	Specifies all routes with a given or larger mask length.
<i>netMaskLen</i>	Specifies a IPv4 or IPv6 subnet mask length (number of bits).
<i>networkPrefixFilter</i>	Specifies an IPv4 or IPv6 address and netmask.
exact	Specifies an exact match with the IP address and subnet mask.

Default

If no address family is specified, IPv4 unicast is the default.

Usage Guidelines

show bgp neighbor now supports v6 unicast and multicast and vpnv4 address families. This command applies to the current VR or VRF context.



Note

If this command displays Bad Source Address, the BGP neighbor IP address is unavailable. Possible causes for this condition include a deleted or unconfigured [VLAN](#) or IP address.

The option `network any / netMaskLen` displays all BGP routes whose mask length is equal to or greater than `maskLength`, irrespective of their network address.

The option `network any / netMaskLen exact` displays all BGP routes whose mask length is exactly equal to `maskLength`, irrespective of their network address.

If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.



Note

For an IPv6 peer, an IPv6 address family must be specified, because an IPv6 peer does not support IPv4 address families. If no address family is specified for an IPv6 peer, the default address-family, i.e. IPv4 unicast is assumed and hence no address-family information appears. Similarly an IPv4 peer only supports IPv4 address families and no address-family information appears if an IPv6 address family is specified.

To display Layer 3 VPN information, you must enter this command in the context of on the [MPLS](#)-enabled VR; it is not supported for BGP neighbors on the CE (VRF) side of the PE router.

Example

The following command displays sample output when an error message is generated when the user attempts to create more than the limit of 2 peering sessions with a Base license on VR-Default (including any child VRFs of VR-Default):

```
# show bgp neighbor
```

Peer	AS	Weight	State	InMsgs
OutMsgs (InQ) Up/Down				
Id-- 10.1.1.1	10	1	IDLE	0 0
(0) 0:0:13:53				
Id-- 20.1.1.1	10	1	IDLE	0 0
(0) 0:0:00:07				

Flags: (d) disabled, (e) enabled, (E) external peer, (I) internal peer

```
switch-model-EXOS.24 # create bgp neighbor 30.1.1.1 remote-AS-number 10
Error: Core license is required to configure more than 2 peers

Base License (User VR):

* (vr user1) switch-model-EXOS.29 # create bgp neighbor 30.1.1.1 remote-AS-number 50
Error: vr user1: Core license is required to configure peer on VR other than VR-Default
```

The following command displays sample output for **show bgp neighbor summary**:

```
# show bgp neighbor

Peer                AS      Weight      State          InMsgs OutMsgs(InQ)  Up/Down
-----
Ie-- 11.0.0.2       100    0           OPENSENT       0      9      (0)  ) 0:8:27:21
Ie-- 3001::1        100    0           ESTABLISEHD    4      3      (0)  ) 0:8:27:21

Flags:  (d) disabled, (e) enabled, (E) external peer, (I) internal peer
        (m) EBGP multihop, (r) route reflector client

BGP Peer Statistics
Total Peers      : 2
EBGP Peers       : 0
IBGP Peers       : 2
RR Client        : 0
EBGP Multihop    : 0
Enabled          : 2
Disabled         : 0
```

The following example displays show output for an IPv4 peer:

```
# show bgp neighbor 192.168.66.2

Peer Description      :
EBGP Peer             : 192.168.66.2      AS              : 38
Enabled               : Yes               OperStatus      : Up
Weight                : 1                 Shutdown-Priority : 1024
ConnectRetry          : 120               MinAsOrig        : 30
HoldTimeCfg           : 180               KeepaliveCfg      : 60
Source Interface      : Not configured    RRClient         : No
EBGP-Multihop         : No                 Remove Private AS : No
BFD                   : Off                BFD Status       : Not Required
Capabilities Config   : ipv4-unicast,ipv4-multicast,4-Byte-As,route-refresh (old &
new),l2vpn-evpn
Policy for NLRI Type  ipv4-unicast
  In Policy           : None
  Out Policy          : None
  NextHopSelf         : Disabled           Send Communities   : No
  Soft Input Recfg    : Disabled           Allow Looped AS-Path: No
  NextHopUnchanged    : Disabled
.
.
.
Policy for NLRI Type  ipv4-vxlan
  In Policy           : None
  Out Policy          : None
  NextHopSelf         : Enabled           Send Communities   : No
  Soft Input Recfg    : Disabled           Allow Looped AS-Path: No
  NextHopUnchanged    : Disabled
Policy for NLRI Type  l2vpn-evpn
  In Policy           : None
  Out Policy          : None
  NextHopSelf         : Disabled           Send Communities   : No
  Soft Input Recfg    : Disabled           Allow Looped AS-Path: No
  NextHopUnchanged    : Enabled
State                 : ESTABLISHED
FSM Up since          : Sat May 5 04:05:30 2018
(Duration: 0:0:08:19)
Remote Addr           : 192.168.66.2      Local Addr        : 192.168.66.1
Remote Port           : 179               Local Port         : 51612
Remote RouterId       : 1.0.0.38          Local RouterId     : 1.0.0.25
```

```

HoldTimeNegotiated : 180                      KeepAliveNegotiated : 60
FsmTransitions      : 1
InUpdateElapsedTime : 00:00:08:25             InMsgElapsedTime      : 0:0:08:25
InUpdates           : 2                       OutUpdates (in TxQ)   : 3 (0)
InTotalMsgs         : 14                     OutTotalMsgs          : 15
InRouteRefreshes    : 0                     OutRouteRefreshes     : 0
Route Statistics for NLRI Type ipv4-unicast
  Received          : 1                       Accepted              : 1
  Rejected          : 0                       Active                 : 1
  Suppressed        : 0
.
.
.
Route Statistics for NLRI Type l2vpn-evpn
  Received          : 0                       Accepted              : 0
  Rejected          : 0                       Active                 : 0
  Suppressed        : 0
Capabilities Tx      : ipv4-unicast,ipv4-multicast,4-Byte-As,route-refresh (old &
new),l2vpn-evpn
Capabilities Rx      : ipv4-unicast,ipv4-multicast,4-Byte-As,route-refresh (old &
new),l2vpn-evpn
NLRI for the session: ipv4-unicast,ipv4-multicast,ipv4-vxlan,l2vpn-evpn
Last State           : ESTABLISHED             Last Event             : RX_KEEP
LastError            : 'Cease - Peer Connection Rejected' (RX) on: Sat May 5 04:05:15 2018

BGP Peer Statistics
  Total Peers        : 1
  EBGp Peers         : 1                       IBGP Peers             : 0
  RR Client          : 0                       EBGp Multihop          : 0
  Enabled            : 1                       Disabled               : 0

```

The following example displays output for an IPv6 peer:

```

# show bgp neighbor 3001::1 detail
EBGP Peer      : 3001::1                      AS                     : 5
Enabled        : Yes                          OperStatus             : Up
Weight         : 1                           Shutdown-Priority       : 1024
ConnectRetry   : 120                        MinAsOrig               : 15
HoldTimeCfg    : 180                        KeepaliveCfg            : 60
Source Interface : Not configured             RRClient                : No
EBGP-Multihop  : No                          Remove Private AS       : No
BFD            : Off                          BFD Status              : Inactive
Capabilities Config : ipv6-unicast, ipv6-multicast, 4-Byte-As, route-refresh
Policy for NLRI Type ipv6-unicast
  In Policy      : None
  Out Policy     : None
  NextHopSelf    : Disabled                    Send Communities       : No
  Soft Input Recfg : Disabled                  Allow Looped AS-Path   : No
  RFD HalfLife   : 0m                          RFD Reuse               : 0
  RFD Suppress   : 0                          RFD Max-Suppress       : 0m
Policy for NLRI Type ipv6-multicast
  In Policy      : None
  Out Policy     : None
  NextHopSelf    : Disabled                    Send Communities       : No
  Soft Input Recfg : Disabled                  Allow Looped AS-Path   : No
  RFD HalfLife   : 0m                          RFD Reuse               : 0
  RFD Suppress   : 0                          RFD Max-Suppress       : 0m
State           : ESTABLISHED
FSM Up since    : Mon Apr 19 21:20:02 2010 (Duration: 0:0:00:23)
Remote Addr     : 3001::1                      Local Addr              : 3001::6
Remote Port     : 56539                          Local Port               : 179
Remote RouterId : 1.0.0.5                        Local RouterId           : 1.0.0.6
HoldTimeNegotiated : 180                      KeepAliveNegotiated     : 60
FsmTransitions  : 5

```

```

InUpdateElapsedTime : 00:00:00:23      InMsgElapsedTime      : 0:0:00:23
InUpdates           : 1                  OutUpdates (in TxQ)   : 4 (0)
InTotalMsgs         : 1                  OutTotalMsgs          : 4
InRouteRefreshes    : 0                  OutRouteRefreshes     : 0
Route Statistics for NLRI Type ipv6-unicast
    Received         : 6
    Accepted          : 6
    Rejected          : 0
    Active            : 6
    Suppressed        : 0
Route Statistics for NLRI Type ipv6-multicast
    Received         : 0
    Accepted          : 0
    Rejected          : 0
    Active            : 0
    Suppressed        : 0
Capabilities Tx      : ipv6-unicast, ipv6-multicast, 4-Byte-AS, route-refresh (old &
new), vpnv4
Capabilities Rx      : ipv6-unicast, ipv6-multicast
NLRI for the session: ipv6-unicast, ipv6-multicast
Error                : 'Hold Timer Expired'          Tx: 3      Rx: 0
Last State           : ESTABLISHED                   Last Event      : RX_UPDATE
LastError            : 'Hold Timer Expired' (TX) on: Mon Apr 19 20:50:26 2010
BGP Peer Statistics
    Total Peers      : 1
    EBGPeers         : 1                  IBGP Peers      : 0
    RR Client        : 0                  EBGPeers        : 0
    Enabled           : 1
    Disabled          : 0

```

The following example displays show output for transmitted routes:

```

#show bgp 11.0.0.2 transmitted-routes all
Advertised Routes:
Destination      LPref Weight MED      Peer      Next-Hop      AS-Path
-----
>? 1.1.1.1/32    100
>? 11.0.0.0/24   100
>? 101.0.0.0/24  100
>? 103.0.0.0/24  100
>? 103.0.0.1/32  100
Flags: (*) Preferred BGP route, (>) Active, (d) Suppressed, (h) History
       (s) Stale, (m) Multipath, (u) Unfeasible
Origin: (?) Incomplete, (e) EGP, (i) IGP
BGP Route Statistics
    Advertised Routes : 5

```

The following example displays show output for rejected routes:

```

# show bgp 11.0.0.2 rejected-routes all
Rejected Routes:
Destination      LPref Weight MED      Peer      Next-Hop      AS-Path
-----
u ? 1.1.1.1/32    100
Flags: (*) Preferred BGP route, (>) Active, (d) Suppressed, (h) History
       (s) Stale, (m) Multipath, (u) Unfeasible
Origin: (?) Incomplete, (e) EGP, (i) IGP
BGP Route Statistics
    Total Rxed Routes : 5
    Rejected Routes    : 1
    Unfeasible Routes  : 1

```

The following example displays show output for accepted routes:

```
# show bgp 11.0.0.2 accepted-routes all
Rejected Routes:
Destination  LPref Weight MED      Peer      Next-Hop      AS-Path
-----
>? 11.0.0.0/24      100
>? 101.0.0.0/24     100
>? 103.0.0.0/24     100
>? 103.0.0.1/32     100
11.0.0.1      100
11.0.0.1      100
11.0.0.1      100
11.0.0.1      100

Flags: (*) Preferred BGP route, (>) Active, (d) Suppressed, (h) History
       (s) Stale, (m) Multipath, (u) Unfeasible
Origin: (?) Incomplete, (e) EGP, (i) IGP
BGP Route Statistics
  Total Rxd Routes : 5
  Feasible Routes  : 4
  Active Routes    : 4
```

The following example shows BGP IPv4 VXLAN routes received from the BGP neighbor located at 192.168.68.1:

```
# show bgp neighbor 192.168.68.1 ipv4-vxlan received-routes all

Routes:

      LTEP      VNI      Peer      Next-Hop      LPref Weight MED      AS-Path
-----
i  1.0.0.15/32      777      192.168.68.1 192.168.68.1      100   1      0   15
```

The following example shows L2VPN EVPN accepted and rejected MAC/IP routes for neighbor 192.168.120.119:

```
# show bgp neighbor 192.168.120.119 l2vpn-evpn mac-ip received-routes all

EVPN MAC/IP Routes:

      MAC      IP      VNI      Peer      Next-
Hop      AS-Path
-----
-
*>? RD: 65535:4294967295      ESI: 00:01:02:03:04:05:06:07:08:09
      00:04:96:98:87:62      192.168.110.109      16777215      192.168.120.119
192.168.120.119      50 40 30 20 10

*>i  RD: 4294967295:65535      ESI: 00:01:02:03:04:05:06:07:08:09
      00:04:96:98:87:64      192.168.110.110      777      192.168.120.119
192.168.120.119      50 40 30 20 10
```

The following example shows all Type 5 received routes to the neighbor at 192.168.99.2:

```
# show bgp neighbor 192.168.99.2 l2vpn-evpn ip-prefix received-routes all

Routes:
EVPN IP-Prefix Routes:
      IP Prefix/Length      VNI
Peer      Next-Hop      AS-Path
-----
-----
```

```

?  RD: 1.0.0.25:48      ESI: 00:00:00:00:00:00:00:00:00:00
    125.126.127.128/32      777
192.168.99.2            1.0.0.25      38 25
?  RD: 1.0.0.25:49      ESI: 00:00:00:00:00:00:00:00:00:00
    1.0.0.0/8              778
192.168.99.2            1.0.0.25      38 25

```

History

This command was first available in ExtremeXOS 10.1.

The any / *netMaskLen* options were added in ExtremeXOS 11.0.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Support for Layer 3 VPNs was added in ExtremeXOS 15.3.

Support for BFD on neighbors was added in ExtremeXOS 21.1

Support for IPv4 VXLAN was added in ExtremeXOS 22.3.

Support for L2VPN EVPN was added in ExtremeXOS 30.2.

Support for EVPN auto-discovery and EVPN Ethernet Segment routes was added in ExtremeXOS 30.4.

Support for Type 5 routes (**ip-prefix**) was added in ExtremeXOS 30.6.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

show bgp peer-group

```
show bgp peer-group {detail | peer-group-name {detail}}
```

Description

Displays the peer groups configured in the system.

Syntax Description

detail	Specifies to display the information in detailed format.
<i>peer-group-name</i>	Specifies a peer group.

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, are displayed.

If no peer group name is specified, all the peer group information is displayed.

This command applies to the current VR or VRF context.

Example

The following command displays information for the outer peer group:

```
* (debug) Summit-PC.19 # show bgp peer-group "outer"
Peer Group           : outer
Enabled              : No                AS                : 65551
Router Enabled       : Yes                Weight             : 1
ConnectRetry        : 120                MinAsOrig           : 15
HoldTimeCfg         : 180                KeepaliveCfg        : 60
Source Interface     : Not configured    RRClient            : No
Remove Private AS   : No                Router-Alert        : Disabled
Capabilities Config  : ipv4-unicast ipv4-multicast route-refresh 4-Byte-AS
Policy for NLRI Type ipv4-unicast
In Policy           : None
Out Policy          : None
NextHopSelf         : Disabled            Send Communities    : No
Soft Input Recfg    : Disabled            Allow Looped AS-Path: No
Policy for NLRI Type ipv4-multicast
In Policy           : None
Out Policy          : None
NextHopSelf         : Disabled            Send Communities    : No
Soft Input Recfg    : Disabled            Allow Looped AS-Path: No
Peers               : 11.11.11.11
BGP Peer Group Statistics
Total Peer Groups   : 1
Enabled             : 0
Disabled            : 1
```

The following command displays information for the pg2 peer group with next-hop unchanged disabled:

```
Peer Group           : pg2
Enabled              : No                AS                : Not configured
Router Enabled       : Yes                Weight             : 1
ConnectRetry        : 120                MinAsOrig           : 15
HoldTimeCfg         : 180                KeepaliveCfg        : 60
Source Interface     : Not configured    RRClient            : No
Remove Private AS   : No
Capabilities Config  : ipv4-unicast,ipv4-multicast,route-refresh (old & new)
Policy for NLRI Type ipv4-unicast
In Policy           : None
Out Policy          : None
NextHopSelf         : Disabled            Send Communities    : No
Soft Input Recfg    : Disabled            Allow Looped AS-Path: No
NextHopUnchanged    : Disabled
Policy for NLRI Type ipv4-multicast
In Policy           : None
Out Policy          : None
NextHopSelf         : Disabled            Send Communities    : No
Soft Input Recfg    : Disabled            Allow Looped AS-Path: No
NextHopUnchanged    : Disabled
```



```

Policy for NLRI Type vpnv4
  In Policy      : None
  Out Policy     : None
  NextHopSelf    : Disabled
  Soft Input Recfg : Disabled
  NextHopUnchanged : Disabled
  Send Communities : Standard, Extended
  Allow Looped AS-Path: No
Policy for NLRI Type ipv6-unicast
  In Policy      : None
  Out Policy     : None
  NextHopSelf    : Disabled
  Soft Input Recfg : Disabled
  NextHopUnchanged : Disabled
  Send Communities : No
  Allow Looped AS-Path: No
Policy for NLRI Type ipv6-multicast
  In Policy      : None
  Out Policy     : None
  NextHopSelf    : Disabled
  Soft Input Recfg : Disabled
  NextHopUnchanged : Disabled
  Send Communities : No
  Allow Looped AS-Path: No
Policy for NLRI Type ipv4-vxlan
  In Policy      : None
  Out Policy     : None
  NextHopSelf    : Disabled
  Soft Input Recfg : Disabled
  NextHopUnchanged : Disabled
  Send Communities : No
  Allow Looped AS-Path: No
Policy for NLRI Type l2vpn-evpn
  In Policy      : None
  Out Policy     : None
  NextHopSelf    : Disabled
  Soft Input Recfg : Disabled
  NextHopUnchanged : Disabled
  Send Communities : Standard, Extended
  Allow Looped AS-Path: No
Peers              : 192.168.48.30

```

The following command displays information for the pg2 peer group with next-hop unchanged enabled:

```

Peer Group      : pg2
Enabled         : No
Router Enabled  : Yes
ConnectRetry    : 120
HoldTimeCfg     : 180
Source Interface : Not configured
Remove Private AS : No
AS              : Not configured
Weight         : 1
MinAsOrig      : 15
KeepaliveCfg    : 60
RRClient       : No
Capabilities Config : ipv4-unicast,ipv4-multicast,route-refresh (old & new)
Policy for NLRI Type ipv4-unicast
  In Policy      : None
  Out Policy     : None
  NextHopSelf    : Disabled
  Soft Input Recfg : Disabled
  NextHopUnchanged : Disabled
  Send Communities : No
  Allow Looped AS-Path: No
Policy for NLRI Type ipv4-multicast
  In Policy      : None
  Out Policy     : None
  NextHopSelf    : Disabled
  Soft Input Recfg : Disabled
  NextHopUnchanged : Disabled
  Send Communities : No
  Allow Looped AS-Path: No
Policy for NLRI Type vpnv4
  In Policy      : None
  Out Policy     : None
  NextHopSelf    : Disabled
  Soft Input Recfg : Disabled
  NextHopUnchanged : Disabled
  Send Communities : Standard, Extended
  Allow Looped AS-Path: No
Policy for NLRI Type ipv6-unicast
  In Policy      : None
  Out Policy     : None
  NextHopSelf    : Disabled
  Send Communities : No

```

```

Soft Input Recfg : Disabled      Allow Looped AS-Path: No
NextHopUnchanged : Disabled
Policy for NLRI Type ipv6-multicast
In Policy       : None
Out Policy      : None
NextHopSelf     : Disabled      Send Communities   : No
Soft Input Recfg : Disabled      Allow Looped AS-Path: No
NextHopUnchanged : Disabled
Policy for NLRI Type ipv4-vxlan
In Policy       : None
Out Policy      : None
NextHopSelf     : Disabled      Send Communities   : No
Soft Input Recfg : Disabled      Allow Looped AS-Path: No
NextHopUnchanged : Disabled
Policy for NLRI Type l2vpn-evpn
In Policy       : None
Out Policy      : None
NextHopSelf     : Disabled      Send Communities   : Standard, Extended
Soft Input Recfg : Disabled      Allow Looped AS-Path: No
NextHopUnchanged : Enabled
Peers           : 192.168.48.30

```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the *BGP* feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

show bgp routes summary

```
show bgp routes {address-family [ipv4-unicast | ipv4-multicast | ipv6-unicast | ipv6-multicast | vpn4 | ipv4-vxlan]} summary {vr vr_name}
```

Description

Displays a summary of the *BGP* route information base (RIB).

Syntax Description

ipv4-unicast	Specifies the IPv4 unicast address family.
ipv4-multicast	Specifies an IPv4 multicast address family.
ipv6-unicast	Specifies the IPv6 unicast address family.
ipv6-multicast	Specifies an IPv6 multicast address family.
vpn4	Specifies the VPNv4 address family for Layer 3 VPN support.
ipv4-vxlan	Specifies an IPv4 VXLAN address family.

vr	Specifies viewing routes associated with a specified virtual router.
<i>vr_name</i>	Specifies the VR name. If not specified, the VR of current command context is used.

Default

If no address family is specified, IPv4 unicast is the default.

If not specified, the VR of current command context is used.

Usage Guidelines

If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.

To display Layer 3 VPN information, you must enter this command in the context of on the *MPLS*-enabled VR; it is not supported for BGP neighbors on the CE (VRF) side of the PE router.

When the `show bgp routes summary` command is issued with **address-family vpnv4**, the command will impact the behavior of PE to PE neighbor sessions and display/clear the VPN-IPv4 RIB of BGP.

Example

The following command displays a summary of the BGP route information base (RIB) for IPv4 multicast:

```
show bgp routes address-family ipv4-multicast summary
```

The following example displays VPN routes with RD 100:1:

```
virtual-router corpl_vrf
  show bgp routes address-family vpnv4 rd 100:1
```

History

This command was first available in ExtremeXOS 10.1.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Support for Layer 3 VPNs was added in ExtremeXOS 15.3.

Support for IPv4 VXLAN was added in ExtremeXOS 22.3.

Support for user-specified VRs was added in ExtremeXOS 30.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

show bgp routes

For IPv4 and IPv6 address families:

```
show bgp routes {address-family [ipv4-unicast | ipv4-multicast | ipv6-unicast | ipv6-multicast | ipv4-vxlan | {l2vpn-evpn [inclusive-multicast | mac-ip | auto-discovery | esi | ip-prefix]}]} {detail} [ipv4-vxlan | all | as-path path-expression | community [no-advertise | no-export | no-export-subconfed | number community_number | autonomous-system-idbgp-community] | network [any/netMaskLen | networkPrefixFilter] {exact}] {vr vr_name}
```

For the VPNv4 address family:

```
show bgp routes address-family vpn4 {detail} [all | as-path path-expression | community [no-advertise | no-export | no-export-subconfed | number community-number | autonomous-system-idbgp-community] | rd rd | network [any/netMaskLen | networkPrefixFilter] {exact}] {vr vr_name}
```

Description

Displays the BGP route information base (RIB).

Syntax Description

ipv4-unicast	Specifies the IPv4 unicast address family.
ipv4-multicast	Specifies an IPv4 multicast address family.
ipv6-unicast	Specifies the IPv6 unicast address family.
ipv6-multicast	Specifies an IPv6 multicast address family.
ipv4-vxlan	Specifies an IPv4 VXLAN address family.
l2vpn-evpn	Specifies an L2VPN EVPN address family.
inclusive-multicast	Displays EVPN inclusive-multicast (type 3) routes
mac-ip	Displays EVPN MAC/IP (type 2) routes.
auto-discovery	Displays EVPN auto-discovery (type 1) routes.
esi	Displays EVPN Ethernet segment (type 4) routes.
ip-prefix	Displays EVPN ip-prefix (type 5) routes.
vpn4	Specifies the VPNv4 address family for Layer 3 VPN support.
all	Specifies all routes.
<i>path-expression</i>	Displays routes that match the specified AA path expression.

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.
<i>community_number</i>	Specifies a community number.
<i>autonomous-system-id</i>	Specifies an autonomous system ID (0-65535).
<i>bgp-community</i>	Specifies the BGP community number.
rd	Specifies the Route Distinguisher (RD) value for the Layer 3 VPN routes for which you want to clear flap statistics.
any	Specifies all routes with a given or larger mask length.
<i>netMaskLen</i>	Specifies a IPv4 or IPv6 subnet mask length (number of bits).
<i>networkPrefixFilter</i>	Specifies an IPv4 or IPv6 address and netmask.
exact	Specifies an exact match with the IP address and subnet mask.
vr	Specifies viewing routes associated with a specified virtual router.
<i>vr_name</i>	Specifies the VR name. If not specified, the VR of current command context is used.

Default

If no address family is specified, IPv4 unicast is the default.

If not specified, the VR of current command context is used.

Usage Guidelines

The option network any / *netMaskLen* displays all BGP routes whose mask length is equal to or greater than *maskLength*, irrespective of their network address.

The option network any / *netMaskLen* exact displays all BGP routes whose mask length is exactly equal to *maskLength*, irrespective of their network address.

To display Layer 3 VPN information, you must enter this command in the context of on the [MPLS](#)-enabled VR; it is not supported for BGP neighbors on the CE (VRF) side of the PE router.

You can only execute the show for **vpn4** address family in a VR context. If you execute this command in a VRF context, the “Cannot execute command in VRF context” error is displayed.

This command applies to the current VR or VRF context.

If you do not specify an address family, this command applies to the IPv4 unicast address family. To apply this command to an address family other than the IPv4 unicast address family, you must specify the address family.

Example

The following command displays detailed information about all BGP routes:

```
* Switch.5 # show bgp routes all
Received Routes:
Destination          LPref Weight MED      Peer          Next-Hop      AS-Path
-----
*>? 1.1.1.1/32        100    0                11.0.0.1      11.0.0.1      100
* ? 11.0.0.0/24       100    0                11.0.0.1      11.0.0.1      100
*>? 101.0.0.0/24      100    0                11.0.0.1      11.0.0.1      100
u ? 103.0.0.0/24      100    0                11.0.0.1      11.0.0.1      100
Flags: (*) Preferred BGP route, (>) Active, (d) Suppressed, (h) History
       (s) Stale, (m) Multipath, (u) Unfeasible

Origin: (?) Incomplete, (e) EGP, (i) IGP

BGP Route Statistics
  Total Rxed Routes : 4
  Feasible Routes   : 3
  Active Routes     : 2
  Rejected Routes   : 0
  Unfeasible Routes : 1
Route Statistics on Session Type
  Routes from Int Peer: 4
  Routes from Ext Peer: 0
```

The following example displays a detailed show output:

```
Route: 11.0.0.0/24, Peer 11.0.0.1, Unfeasible
Origin Incomplete, Next-Hop 11.0.0.1, LPref 100, MED 0
Weight 0, RR Orig ID 0.0.0.0
AS-Path: 100
DampInfo: Penalty 0 Flapped 0 times in 00:10:47
Flags: (*) Preferred BGP route, (>) Active, (d) Suppressed, (h) History
       (s) Stale, (m) Multipath, (u) Unfeasible
Origin: (?) Incomplete, (e) EGP, (i) IGP
BGP Route Statistics
  Total Rxed Routes : 1
  Feasible Routes   : 0
  Active Routes     : 0
  Rejected Routes   : 5
  Unfeasible Routes : 1
Route Statistics on Session Type
  Routes from Int Peer : 5
  Routes from Ext Peer : 0
```

The following command displays BGP information for the IPv6 address family:

```
Switch.21 # show bgp routes address-family ipv6-unicast all
Received Routes:
Destination          LPref Weight MED      Peer          Next-Hop      AS-Path
-----
*>? 2001::/64         100    0                3001::1       3001::1       120
    3000:::1          100    100              3001::1       3001::1       200
Flags: (*) Preferred BGP route, (>) Active, (d) Suppressed, (h) History
       (s) Stale, (m) Multipath, (u) Unfeasible

Origin: (?)
       Incomplete, (e) EGP, (i) IGP

BGP Route Statistics
```

```

Total Rxed Routes : 1
Feasible Routes   : 1
Active Routes     : 1
Rejected Routes   : 0
Unfeasible Routes : 0

Route Statistics on Session Type
  Routes from Int Peer: 1
    Routes from Ext Peer: 0

```

The following example displays detailed show output for the IPv6 address family:

```

switch.21 # show bgp routes address-family ipv6-unicast all
Route: 2001::/64, Peer 3000::1,
Unfeasible, Origin Incomplete,
Next-Hop 3001::1,
LPref 100, MED 0,
Weight 0, RR Orig ID 0.0.0.0
AS-Path: 100
DampInfo: Penalty 0 Flapped 0 times in 00:10:47

Route: 2002::/64, Peer 3000::1,
Active, Origin Incomplete,
Next-Hop 3001::1,
LPref 100, MED 0,
Weight 0, RR Orig ID 0.0.0.0
AS-Path: 100
DampInfo: Penalty 0 Flapped 0 times in 00:10:47

BGP Route Statistics
  Total Rxed Routes : 2
  Feasible Routes   : 1
  Active Routes     : 1
  Rejected Routes   : 0
  Unfeasible Routes : 1
Route Statistics on Session Type
  Routes from Int Peer: 2
    Routes from Ext Peer: 0

```

The following examples display detailed show output for the IPv4 address family:

```

switch.21 # show bgp routes address-family vpnv4 all
Received Routes:
Destination
Peer                Next-Hop                LPref Weight  MED
AS-Path
-----
*>? 100:1:10.0.0.0/8      100    0      120
    11.0.0.2              11.0.0.2
    100, 200
Flags: (*) Preferred BGP route, (>) Active, (d) Suppressed, (h) History
      (s) Stale, (m) Multipath, (u) Unfeasible

Origin: (?) Incomplete, (e) EGP, (i) IGP

BGP Route Statistics
  Total Rxed Routes : 1
  Feasible Routes   : 1
  Active Routes     : 1
  Rejected Routes   : 0
  Unfeasible Routes : 0
Route Statistics on Session Type
  Routes from Int Peer: 1
    Routes from Ext Peer: 0

```

```

-----
switch.21 # show bgp routes address-family ipv6-unicast all
Route: 100:1:10.0.0.0/8, Peer 11.0.0.2,
Unfeasible, Origin Incomplete,
Next-Hop 11.0.0.2,
LPref 100, MED 0,
Weight 0, RR Orig ID 0.0.0.0
AS-Path: 100
DampInfo: Penalty 0 Flapped 0 times in 00:10:47

BGP Route Statistics
  Total Rxed Routes : 1
  Feasible Routes   : 0
  Active Routes     : 0
  Rejected Routes   : 0
  Unfeasible Routes : 1
Route Statistics on Session Type
  Routes from Int Peer: 1
  Routes from Ext Peer: 0

```

The following example displays detailed show output for the IPv4 VXLAN address family:

```

# show bgp routes ipv4-vxlan all

Routes:
  LTEP          VNI      Peer      Next-Hop      LPref Weight MED      AS-
Path
-----
* i  1.0.0.15/32      777      192.168.68.1 192.168.68.1   100   1      0
15
Flags: (*) Preferred BGP route, (>) Active, (d) Suppressed, (h) History
       (s) Stale, (m) Multipath, (u) Unfeasible

```

The following example displays L2VPN EVPN MAC/IP routes:

```

# show bgp routes l2vpn-evpn mac-ip all

EVPN MAC/IP Routes:

Hop      MAC          IP          VNI      Peer      Next-
      AS-Path
-----
*>? RD: 65535:4294967295      ESI: 00:01:02:03:04:05:06:07:08:09
    00:04:96:98:87:62      192.168.110.109      16777215      192.168.120.119
192.168.120.119      50 40 30 20 10

*>i RD: 192.168.101.100:65535 ESI: 00:00:00:00:00:00:00:00:00:00
    00:04:96:98:87:63      0.0.0.0      16777215      192.168.0.3
192.168.0.3      50

*>i RD: 4294967295:65535      ESI: 00:01:02:03:04:05:06:07:08:09
    00:04:96:98:87:64      192.168.110.110      777      192.168.120.119
192.168.120.119      50 40 30 20 10

Flags: (*) Preferred BGP route, (>) Active, (d) Suppressed, (h) History
       (s) Stale, (m) Multipath, (u) Unfeasible

Origin: (?) Incomplete, (e) EGP, (i) IGP

```


The following example displays L2VPN EVPN inclusive multicast routes:

```
# show bgp routes l2vpn-evpn inclusive-multicast all

EVPN Inclusive Multicast Routes:

      Originating IP          Tunnel ID          VNI
Peer      Next-Hop          AS-Path
-----
-
*>?  RD: 65535:4294967295   Ethernet Tag: 0   Tunnel Type: 6 (Ingress Replication)
      192.168.110.109       192.168.110.109   16777215
192.168.120.119   192.168.120.119   50 40 30 20 10

Flags: (*) Preferred BGP route, (>) Active, (d) Suppressed, (h) History
      (s) Stale, (m) Multipath, (u) Unfeasible

Origin: (?) Incomplete, (e) EGP, (i) IGP
```

The following example shows all Type 5 routes:

```
# show bgp routes l2vpn-evpn ip-prefix all

Routes:
EVPN IP-Prefix Routes:
      IP Prefix/Length          VNI
Peer      Next-Hop          AS-Path
-----
-
?  RD: 1.0.0.25:48           ESI: 00:00:00:00:00:00:00:00:00:00
      125.126.127.128/32       777
192.168.99.2           1.0.0.25           38 25
?  RD: 1.0.0.25:49           ESI: 00:00:00:00:00:00:00:00:00:00
      1.0.0.0/8               778
192.168.99.2           1.0.0.25           38 25

Flags: (*) Preferred BGP route, (>) Active, (d) Suppressed, (h) History
      (s) Stale, (m) Multipath, (u) Unfeasible

Origin: (?) Incomplete, (e) EGP, (i) IGP

BGP Route Statistics EVPN route type 5
Total Rxed Routes : 2
Feasible Routes   : 2
Active Routes     : 0
Rejected Routes   : 0
Unfeasible Routes : 0
Route Statistics on Session Type
Routes from Int Peer: 0
Routes from Ext Peer: 2
```

History

This command was first available in ExtremeXOS 10.1.

The any / *netMaskLen* options were added in ExtremeXOS 11.0.

This command required a specific license in ExtremeXOS 11.1.

Support for IPv6 was added in ExtremeXOS 12.6 BGP.

Support for Layer 3 VPNs was added in ExtremeXOS 15.3.

Support for IPv4 VXLAN was added in ExtremeXOS 22.3.

Support for L2VPN EVPN was added in ExtremeXOS 30.2.

Support for user-specified VRs was added in ExtremeXOS 30.3.

Support for EVPN auto-discovery and EVPN Ethernet Segment routes was added in ExtremeXOS 30.4.

Support for Type 5 routes (**ip-prefix**) was added in ExtremeXOS 30.6.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the BGP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

show bootprelay

```
show bootprelay
```

Description

Displays the *DHCP*/BOOTP Relay statistics and the configuration for the VRs.

Syntax Description

ipv4	Specifies unconfiguring the DHCPv4 BOOTP Relay service (default).
ipv6	Specifies unconfiguring the DHCPv6 BOOTP Relay service.
vlan	Unconfigures BOOTP relay for a specified <i>VLAN</i> .
<i>vlan_name</i>	Specifies the VLAN name.
include-secondary	Removes the include-secondary configuration for the specified VLAN.

Default

None.

Usage Guidelines

The fields displayed in the DHCP Information Option 82 section depend on the configuration defined by the `configure bootprelay dhcp-agent information policy [drop | keep | replace]` command. If the policy configured is keep, the Requests unmodified counter appears. If the policy configured is replace, the Requests replaced counter appears. And if the drop policy is configured, the Requests dropped counter appears.

The Opt82 added to Requests counter indicates the number of DHCP requests to which the BOOTP Relay agent (the switch) has added its own option 82 information.

Example

The following example displays the DHCP/BOOTP relay statistics for existing VRs:

```
Switch.1 # show bootprelay
Bootprelay : Enabled on virtual router "VR-Default"
DHCP Relay Agent Information Option : Enabled on virtual router "VR-Default"
DHCP Relay Agent Information Check : Enabled on virtual router "VR-Default"
DHCP Relay Agent Information Policy : Replace
DHCP Relay Agent Information Remote-ID : "default"
Bootprelay servers for virtual router "VR-Default":
Destination: 10.127.8.1
DHCP/BOOTP relay statistics for virtual router "VR-Default"
Received from client =          2  Received from server =          2
Requests relayed      =          2  Responses relayed      =          2
DHCP Discover         =          1  DHCP Offer          =          1
DHCP Request          =          1  DHCP Ack            =          1
DHCP Decline          =          0  DHCP NACK           =          0
DHCP Release          =          0
DHCP Inform           =          0
DHCP Information Option 82 packets statistics for virtual router "VR-Default"
Received from client  =          0  Received from server =          2
Requests replaced     =          0  Responses dropped    =          0
Opt82 added to Requests =          2
Note: Default Remote-ID : System MAC Address
```

The following example shows DHCP/BOOTP relay that is disabled for the VR, but enabled on some VLANs:

```
# show bootprelay
Bootprelay : Disabled on virtual router "VR-Default", but enabled on some VLANs.
Include Secondary : Disabled
DHCP Relay Agent Information Option : Disabled on virtual router "VR-Default"

DHCP/BOOTP relay statistics for virtual router "VR-Default"
Received from client =          0  Received from server =          0
Requests relayed     =          0  Responses relayed     =          0
DHCP Discover        =          0  DHCP Offer            =          0
DHCP Request         =          0  DHCP Ack              =          0
DHCP Decline         =          0  DHCP NACK             =          0
DHCP Release         =          0
DHCP Inform          =          0
```

History

This command was first available in ExtremeXOS 10.1.

Information about DHCP/BOOTP relay being disabled for the VR, but enabled on some VLANs was added in ExtremeXOS 22.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show bootprelay configuration

```
show bootprelay configuration {ipv4 | ipv6} {{vlan vlan_name } | {vr
vr_name}}
```

Description

Displays the enabled/disabled configuration of BOOTP relay on one or all VLANs for the specified VR.

Syntax Description

bootprelay	BOOTP Relay Configuration
ipv4	DHCPv4 BOOTP Relay service.
ipv6	DHCPv6 BOOTP Relay service.
vlan	VLAN
<i>vlan_name</i>	Specifies a single VLAN for which to display BOOTP relay configuration information.
vr	Use a specific virtual router name.
<i>vr_name</i>	Specifies a single VR for which to display BOOTP relay configuration information.

Default

IPv4.

If a VR is not specified, this command displays the specified VLANs for the current VR context.

Usage Guidelines



Note
ExtremeXOS DHCPv6 supports three options: OPTION_RELAY_MSG (9), OPTION_INTERFACE_ID (18), and OPTION_REMOTE_ID (37).



Note
It is mandatory to configure BOOTP Relay v6 Agents. The packets are not forwarded if the v6 agents are not configured.

Example

The following example displays the BOOTP relay configuration for all VLANs on the VR-Default virtual router:

```
# show bootprelay configuration vr "VR-Default"
BOOTP Relay : Enabled on virtual router "VR-Default"
VLAN
-----
Default                Disabled
client1                Enabled
serv                   Enabled
```

The following example displays the BOOTP relay configuration for all VLANs in the current VR context:

```
# show bootprelay configuration
BOOTP Relay : Enabled on virtual router "VR-Default"
VLAN
-----
Default                Disabled
client1                Enabled
serv                   Enabled
```

-----	-----
Default	Disabled
client1	Disabled
serv	Disabled

The following example displays the BOOTP relay configuration for VLAN client1:

```
# show bootprelay configuration vlan "client1"
BOOTP Relay : Enabled on virtual router "VR-Default"
VLAN
-----
client1
```

The following example displays the BOOTP relay configuration for IPv6:

```
# show bootprelay configuration ipv6
DHCPv6 BOOTP Relay : Enabled on virtual router "VR-Default"
VLAN
-----
Default
```

The following example displays the BOOTP relay configuration for IPv6 when vr is disabled:

```
# show bootprelay configuration ipv6
DHCPv6 BOOTP Relay : Disabled on virtual router "VR-Default"
```

The following example shows DHCP/BOOTP relay configuration that is disabled for the VR, but enabled on some VLANs:

```
# show bootprelay configuration
DHCPv4 BOOTP Relay : Disabled on virtual router "VR-Default", but enabled on some VLANs.
  Include Secondary      : Disabled
  BOOTP Relay Servers   :
  DHCP Relay Agent Information Option: Disabled
  DHCP Relay Agent Information Check : Disabled
  DHCP Relay Agent Information Policy: Replace
  DHCP Relay Agent Source VLAN      : ltep
```

History

This command was first available in ExtremeXOS 12.4.2.

Information about DHCP/BOOTP relay being disabled for the VR, but enabled on some VLANs was added in ExtremeXOS 22.5.

VLAN to use as the source IP address in the BOOTP relay packet information was added in 30.7.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show bootprelay configuration ipv4

```
show bootprelay
```

Description

Displays various BOOTP Relay configuration details.

Syntax Description

ipv4	Specifies unconfiguring the DHCPv4 BOOTP Relay service (default).
ipv6	Specifies unconfiguring the DHCPv6 BOOTP Relay service.
vlan	Unconfigures BOOTP relay for a specified <u>VLAN</u> .
<i>vlan_name</i>	Specifies the VLAN name.
include-secondary	Removes the include-secondary configuration for the specified VLAN.

Default

Usage Guidelines

Use this command to display various bootprelay configuration details.

Example

The following command displays IPv4 bootprelay statistics:

```
# show bootprelay configuration ipv4
DHCPv4 BOOTP Relay   : Enabled on virtual router "VR-Default"
  Include Secondary   : Enabled (parallel)
  BOOTP Relay Servers :
    DHCP Relay Agent Information Option: Enabled
    DHCP Relay Agent Information Check : Disabled
    DHCP Relay Agent Information Policy: Replace
VLAN                  DHCPv4 BOOTP Relay
-----
VLAN "Default":
  BOOTP Relay          : Enabled
  Include Secondary    : Enabled (sequential)
  BOOTP Relay Servers   : 192.168.1.1
  DHCP Relay Agent Information Option: Enabled
  DHCP Relay Agent Information Check : Disabled
  DHCP Relay Agent Information Policy: Replace
VLAN "v1":
  BOOTP Relay          : Enabled
  Include Secondary    : Enabled (sequential)
  BOOTP Relay Servers   : 10.1.1.1
  DHCP Relay Agent Information Option: Enabled
  DHCP Relay Agent Information Check : Disabled
  DHCP Relay Agent Information Policy: Replace
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show bootprelay configuration ipv6

show bootprelay

Description

Displays various BOOTP Relay configuration details.

Syntax Description

ipv4	Specifies unconfiguring the DHCPv4 BOOTP Relay service (default).
ipv6	Specifies unconfiguring the DHCPv6 BOOTP Relay service.
vlan	Unconfigures BOOTP relay for a specified <u>VLAN</u> .
<i>vlan_name</i>	Specifies the VLAN name.
include-secondary	Removes the include-secondary configuration for the specified VLAN.

Default

N/A.

Usage Guidelines

Use this command to display BOOTP Relay details for IPv6.

Example

The following command :

```
# show bootprelay configuration ipv6
DHCPv6 BOOTP Relay : Enabled on virtual router "VR-Default"
Include Secondary : Enabled (parallel)

VLAN                               DHCPv6 BOOTP Relay
-----
VLAN "Default":
    BOOTP Relay                   : Disabled
VLAN "v1":
    BOOTP Relay                   : Enabled
    Include Secondary             : Enabled (sequential)
VLAN "v2":
    BOOTP Relay                   : Enabled
    Include Secondary             : Enabled (parallel)
```

The following example shows DHCPv6/BOOT Relay configuration that is disabled on the VR, but enabled on some VLANs:

```
# show bootprelay configuration ipv6
DHCPv6 BOOTP Relay : Disabled on virtual router "VR-Default", but enabled on some VLANs.
Include Secondary : Disabled

VLAN                               DHCPv6 BOOTP Relay
-----
```

```
VLAN "Default":
  BOOTP Relay      : Enabled
```

History

This command was first available in ExtremeXOS 15.4.

Information about DHCP/BOOTP relay being disabled for the VR, but enabled on some VLANs was added in ExtremeXOS 22.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show bootprelay dhcp-agent information circuit-id port-information

```
show bootprelay dhcp-agent information circuit-id port-information ports
all
```

Description

Displays the circuit ID sub-option that identifies the port for an incoming *DHCP* request.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays the circuit ID port_info value for all ports:

```
Switch.12 # show bootprelay dhcp-agent information circuit-id port-information ports all
Port          Circuit-ID Port information string
----          -
1             1001
2             1002
3             extreme1
4             1004
5             1005
6             1006
7             1007
8             1008
9             1009
10            1010
```



```

:
:
11          1011
12          1012
:
:
48          1048
49          1049
50          1050
Note: The full Circuit ID string has the form '<Vlan Info>--<Port Info>'

```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show bootprelay dhcp-agent information circuit-id vlan-information

```
show bootprelay dhcp-agent information circuit-id vlan-information
```

Description

Displays the circuit ID sub-options that identify the VLANs on the switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays the circuit ID vlan_info for all VLANs:

```

Summit # show bootprelay dhcp-agent information circuit-id vlan-information
Vlan          Circuit-ID vlan information string
----          -
Default       1
Mgmt          4095
v1            4094
v2            extreme123
Note: The full Circuit ID string has the form '<Vlan Info>--<Port Info>'

```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show bootprelay ipv6

show bootprelay ipv6

Description

Displays various BOOTP Relay configuration details.

Syntax Description

bootprelay	Shows the BOOTP relay information.
ipv6	DHCPv6 BOOTP Relay Service.

Default

Not applicable.

Usage Guidelines

Use this command to display IPv6 bootp relay information.

Example

The following command displays IPv6 bootprelay information:

```
* switch # show bootprelay ipv6
BOOTP Relay: DHCPv6 BOOTP Relay enabled on virtual router "VR-Default"
Include Secondary : Disabled
    BOOTP Relay Servers :2001::1
                        3001::1
                        4001::1
VLAN "Default"      :
    BOOTP Relay      : Enabled
    Interface ID     : 3999 (Default)
    Remote ID        : 00:04:96:52:08:76 (Default)
    Prefix Snooping  : Disabled
VLAN "v1"           :
    BOOTP Relay      : Enabled
    Interface ID     : Interface-Sring1
    Remote ID        :
* switch #
```

When vr is disabled:

```
* SWITCH # show bootprelay ipv6 configuration
BOOTP Relay: DHCPv6 BOOTP Relay disabled on virtual router "VR-Default"
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show bootprelay ipv6 prefix-delegation snooping

```
show bootprelay ipv6 prefix-delegation snooping {vlan} vlan_name
```

Description

Displays the information about snooped IPv6 prefixes delegated via *DHCP*.

Syntax Description

<i>vlan_name</i>	Specifies the <i>VLAN</i> name.
------------------	---------------------------------

Default

N/A.

Usage Guidelines

None.

Example

The following command displays snooped IPv6 prefixes delegated via DHCP for all VLANs.

```
show bootprelay ipv6 prefix-delegation snooping
```

The following is sample output:

Delegated Prefix Gateway	Interface Valid For
-----	-----
2000::/48	v1
2000:::1	7154 secs
3000::/48	v2
3000:::1	992 secs

```
2800::/56                                v3
2800::1                                10 secs
```

The following command displays snooped IPv6 prefixes delegated via DHCP for VLAN v1:

```
show bootprelay ipv6 prefix-delegation snooping vlan v1
```

The following is sample output:

Delegated Prefix Gateway	Interface Valid For
-----	-----
2000::/48	v1
2000::/1	7074 secs

History

This command was first available in ExtremeXOS 15.7.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show cdp

```
show cdp
```

Description

Displays the interval between advertisements, the hold time and the version of the advertisement.

Syntax Description

There are no arguments or keywords for this command.

Default

N/A.

Usage Guidelines

Use this command to display the interval between advertisements, the hold time and the version of the advertisement.

Example

The following command displays specific information on the CDP feature:

```
# show cdp
1234567890123456789012345678901234567890123456789012345678901234567890
```

```

CDP Transmit time           : 60 seconds
CDP Hold time               : 180 seconds
CDP Device ID               : 00:04:96:8B:C2:CA
CDP Enabled ports           : 1-2, 7
Power Available TLV Enabled ports : 1-2,23
CDP Local management address : VLAN Chicago (2001:db8:85a3::7334)

```

History

This command was first available in ExtremeXOS 15.4.

The output of this command was updated in ExtremeXOS 21.1.

Management IP address information was added in ExtremeXOS 22.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show cdp counters

```
show cdp counters {ports port_list}
```

Description

Displays CDP port counter statistics.

Syntax Description

ports	Displays CDP port statistics.
<i>port_list</i>	Port list of CDP ports.

Default

N/A.

Usage Guidelines

Use this command to display CDP port counter statistics.

Example

The following command displays counter statistic for CPD ports:

```

# show cdp counters ports 10
Port    Tx PDU      Rx PDU      Rx Drop      Rx Drop
Count   Count       Count       Checksum Err  Invalid
-----
10      3729        3729        0             0

```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show cdp neighbor

```
show cdp neighbor {detail}
```

Description

Displays information about neighbors.

Syntax Description

detail	Displays detailed information.
---------------	--------------------------------

Default

N/A.

Usage Guidelines

Use this command to display CDP neighbor information.

History

This command was first available in ExtremeXOS 15.4.

The output of this command was updated in ExtremeXOS 21.1

Management IP address information was added in ExtremeXOS 22.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show cdp ports

```
show cdp ports {port_list} {configuration | detail}
```

Description

Displays information about neighbors in ports.

Syntax Description

<i>port_list</i>	Specifies the port list, separated by a comma.
configuration	CDP local port configuration.
detail	Detailed information.

Default

N/A.

Usage Guidelines

Use this command to display CDP port information.

Example

The following command displays neighbor port information on the CDP feature:

```
# show cdp ports

Neighbor Information-----
Port  Device-Id                Hold time  Remote CDP  Port
ID
                                     Version
-----
1      Eni-Extreme-x440-sw> 149          Version-1   Slot: 1, Port: 1
2      00:04:96:8B:9D:B0    160          Version-2   Slot: 1, Port: 2
7      00:04:96:8B:C1:ED    138          Version-2   Slot: 1, Port: 7
```



Note

">" indicates that the value was truncated to the column size in the output.

The following example displays the COS extended neighbor port information on the CDP feature:

```
Local Port Information
-----
Port      Trust      Cos    VoIP VLAN                VoIP VLAN
Advertise
-----
1          Trusted    0      Default                  Solicited
2          Trusted    0      Default                  Unsolicited
```

The following example shows detailed information for port 10:

```
# show cdp port 10 detail

Neighbor Information
-----
Device ID           : X670G2-48x-4q
Port ID (outgoing port) : Slot: 1, Port: 10
Advertisement Version : 2
Platform            : X670G2-48x-4q
Interface            : 10
Holdtime             : 162

Version              :
ExtremeXOS version 22.4.0.5 xos_22.4 by kosharma
```

```

on Fri Jul 14 12:28:36 IST 2017

Native VLAN      : 1
Duplex           : Full
Trust           : Trusted
SysName         : X670G2-48x-4q

```

History

This command was first available in ExtremeXOS 15.4.

The output of this command was updated in ExtremeXOS 21.1.

The **detail** option was added and the output of this command was updated in ExtremeXOS 22.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show cfm detail

```

show cfm {domain_name {association_name {ports port_list} {[end-point
    [up | down]]}}}} detail

```

Description

Displays the MEP CCM database.

Syntax Description

<i>domain_name</i>	Enter the name of the domain for which you want to display the MEP CCM databases.
<i>association_name</i>	IEEE 802.1ag or ITU-T Y.1731 association name.
<i>port_list</i>	Enter the ports in the domain/association for which you want to display the CCM databases.
up	Enter this to display the CCM database on the UP MEP for the specified MA.
down	Enter this to display the CCM database on the DOWN MEP for the specified MA.

Default

N/A.

Usage Guidelines

If you do not specify any parameters or variables, the system displays information on all CCM databases on the switch.

This command displays the following items of the CCM database:

- The name of the domain and association
- Port number
- MP and type
- MAC address of remote end points
- MEP IDs
- Lifetime for CCM messages from each remote end point
- Actual age of CCM messages



Note

The TTL for the CCM messages from the MP you are working on is 3.5 times the transmission interval.

Example

The following command displays the CCM databases on the switch:

```
show cfm detail
```

The following is sample output from this command:

```
# show cfm detail
Domain/      Port      MP      Remote End-Point      Remote End-Point MEP      Life      Age      Flags
Association  MAC Address      IP Address      ID      time
=====
dnsname
10           1          UE      00:04:96:51:5f:15 0.0.0.0          300      3500    650    DMA
dom1
VSNLMEG1     1          DE      -----            0.0.0.0          300      11      0
SMA
           1          DE      -----            0.0.0.0          400      11      0      SMI
           15         DE      -----            0.0.0.0          300      35      0      SMA
           15         DE      -----            0.0.0.0          400      35      0      SMI
dom1
short_ma_name 1          DE      -----            0.0.0.0          300      11      0      SMA
           1          DE      -----            0.0.0.0          400      11      0      SMI
           15         DE      -----            0.0.0.0          300      3500    0      SMA
           15         DE      -----            0.0.0.0          400      3500    0      SMI
dom2
VSNLMEG1     1          UE      00:04:96:51:5f:15 0.0.0.0          300      3500    750    SUA
00:11:22:33:4
VSNLMEG1     1          UE      00:04:96:51:5f:15 0.0.0.0          300      3500    760    DMA
00:11:22:33:4
short_ma_name 1          UE      00:04:96:51:5f:15 10.10.10.2       300      3500    90      DMA
=====
Maintenance Point: (UE) Up End-Point, (DE) Down End-Point
Flags: (S) Static Entry, (D) Dynamic Entry
       CCM Destination MAC: (U) Unicast, (M) Multicast
       Status: (A) Active, (I) Inactive
NOTE: The Domain and Association names are truncated to 13 characters, Lifetime
and Age are in milliseconds.
=====
Total Number of Dynamic Up RMEP      : 3
Total Number of Dynamic Down RMEP    : 0
```

```
Total Number of Active Static RMEP : 5
Total Number of Inactive Static RMEP : 4
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show cfm groups

```
show cfm groups {group_name}
```

Description

This command displays the details of specified or all groups. The information contains group name, group status, LMEP id, the physical port of the LMEP, RMEP ids, registered clients, domain and association names.

Syntax Description

<i>group_name</i>	Group name, maximum of 31 characters.
-------------------	---------------------------------------

Default

N/A.

Usage Guidelines

Use this command to display the details of specified or all groups. The information contains group name, group status, LMEP id, the physical port of the LMEP, RMEP ids, registered clients, domain and association names.

Example

The following output shows the typical output of this command:

```
# sh cfm groups
Group : eapsCfmGrp1      Status : UP
Local MEP      : 11      port   : 41
Remote MEPs    : 10
Client(s)      : eaps
Domain         : MD1
Association    : MD1v2
Group : eapsCfmGrp2      Status : UP
Local MEP      : 12      port   : 31
Remote MEPs    : 13
Client(s)      : eaps
```

```
Domain       : MD1
Association   : MD1v2
```

The following example shows the output for *ERPS* with Y.1731 CCMs:

```
# show configuration cfm
#
# Module dot1ag configuration.
#
create cfm domain string "dom" md-level 5
configure cfm domain "dom" add association meg "VSNLMEG1" vlan "v1"
configure cfm domain "dom" association "VSNLMEG1" ports 1 add end-point down 100
configure cfm domain "dom" association "VSNLMEG1" ports 15 add end-point down 200
configure cfm domain "dom" association "VSNLMEG1" ports 1 end-point down add group "erps-
g1"
configure cfm domain "dom" association "VSNLMEG1" ports 15 end-point down add group "erps-
g2"
configure cfm group "erps-g1" add rmep 300
configure cfm group "erps-g2" add rmep 400
#
# show configuration "erps"
#
# Module erps configuration.
#
create erps erps_major_1
configure erps erps_major_1 add control vlan v1
configure erps erps_major_1 ring-port east 1
configure erps erps_major_1 ring-port west 15
configure erps erps_major_1 timer wait-to-restore 5000
configure erps erps_major_1 cfm port east add group erps-g1
configure erps erps_major_1 cfm port west add group erps-g2*
#
# show cfm detail
```

Domain/ Flags	Port	MP	Remote End-Point MAC Address	Remote End-Point IP Address	MEP ID	Life time	Age
dom							
VSNLMEG1	1	DE	-----	0.0.0.0	300	3500	0
SMA	15	DE	-----	0.0.0.0	400	3500	0
SMA							

```
=====
Maintenance Point: (UE) Up End-Point, (DE) Down End-Point
Flags: (S) Static Entry, (D) Dynamic Entry
       CCM Destination MAC: (U) Unicast, (M) Multicast
       Status: (A) Active, (I) Inactive
NOTE: The Domain and Association names are truncated to 13 characters, Lifetime
      and Age are in milliseconds.
=====

Total Number of Dynamic Up RMEP      : 0
Total Number of Dynamic Down RMEP    : 0
Total Number of Active Static RMEP   : 2
Total Number of Inactive Static RMEP : 0
# show cfm groups

Group : erps-g1 Status : UP
       Local MEP      : 100    port : 1
       Remote MEPs    : 300
       Client(s)      : erps
       Domain         : dom
```

```

      Association      : VSNLMEG1
Group : erps-g2 Status : UP
      Local MEP       : 200      port   : 15
      Remote MEPs     : 400
      Client(s)       : erps
      Domain          : dom
      Association      : VSNLMEG1
#
# disable ports 1
#
# show cfm detail

Domain/      Port      MP      Remote End-Point      Remote End-Point      MEP      Life
Flags
Association          MAC Address          IP Address          ID      time      Age
=====
dom
VSNLMEG1      1      DE      -----      0.0.0.0      300      3500      0
SMI
      15      DE      -----      0.0.0.0      400      3500      0
SMA
=====
Maintenance Point: (UE) Up End-Point, (DE) Down End-Point
Flags: (S) Static Entry, (D) Dynamic Entry
      CCM Destination MAC: (U) Unicast, (M) Multicast
      Status: (A) Active, (I) Inactive
NOTE: The Domain and Association names are truncated to 13 characters, Lifetime
      and Age are in milliseconds.
=====

Total Number of Dynamic Up RMEP      : 0
Total Number of Dynamic Down RMEP    : 0
Total Number of Active Static RMEP    : 1
Total Number of Inactive Static RMEP  : 1
# show cfm groups

Group : erps-g1 Status : DOWN
      Local MEP       : 100      port   : 1
      Remote MEPs     : 300
      Client(s)       : erps
      Domain          : dom
      Association      : VSNLMEG1
Group : erps-g2 Status : UP
      Local MEP       : 200      port   : 15
      Remote MEPs     : 400
      Client(s)       : erps
      Domain          : dom
      Association      : VSNLMEG1
# enable ports 1
#
# show cfm detail

Domain/      Port      MP      Remote End-Point      Remote End-Point      MEP      Life
Flags
Association          MAC Address          IP Address          ID      time      Age
=====
dom
VSNLMEG1      1      DE      -----      0.0.0.0      300      3500      0
SMA
      15      DE      -----      0.0.0.0      400      3500      0
SMA
=====
Maintenance Point: (UE) Up End-Point, (DE) Down End-Point
Flags: (S) Static Entry, (D) Dynamic Entry
      CCM Destination MAC: (U) Unicast, (M) Multicast

```

```

        Status: (A) Active, (I) Inactive
NOTE: The Domain and Association names are truncated to 13 characters, Lifetime
      and Age are in milliseconds.
=====
Total Number of Dynamic Up RMEP      : 0
Total Number of Dynamic Down RMEP    : 0
Total Number of Active Static RMEP   : 2
Total Number of Inactive Static RMEP : 0
# show cfm groups

Group : erps-g1 Status : UP
      Local MEP       : 100      port   : 1
      Remote MEPs     : 300
      Client(s)       : erps
      Domain          : dom
      Association      : VSNLMEG1
Group : erps-g2 Status : UP
      Local MEP       : 200      port   : 15
      Remote MEPs     : 400
      Client(s)       : erps
      Domain          : dom
      Association      : VSNLMEG1

```

History

This command was first available in ExtremeXOS 15.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show cfm segment frame-delay statistics

```
show cfm segment frame-delay statistics {segment-name} {mep mep_id}
```

Description

This command displays frame-delay information for the given CFM segment.

Syntax Description

<i>segment_name</i>	An alpha numeric string identifying the segment name.
mep	Maintenance association End Point.
<i>mep_id</i>	MEP-ID. The range is 1-8191.

Default

N/A.

Usage Guidelines

Use this command to display the delay for the last received frame, the minimum, maximum and average delay, and the delay variance during the current transmission. When the segment name is not specified, only the segments which have valid statistics alone are displayed. When the segment name is specified, that particular segment's information, although not present, is displayed.

Example

The following command displays the frame delay statistics for the CFM segment:

```
show cfm segment frame-delay statistics
```

Following is sample output for this command:

```

-----
Segment Name      Mep      Recent  Min      Max      Mean      Jitter      Errored
                  ID       Delay   Delay   Delay   Delay   Delay      Frames*
                  (ms)    (ms)    (ms)    (ms)    (ms)    (ms)
-----
segment1          ----    0.000   0.000   0.000   0.000   0.000       0
segment2          100    0.000   0.000   0.000   0.000   0.000       0
                  200    0.000   0.000   0.000   0.000   0.000       0
segment3          100    0.000   0.000   0.000   0.000   0.000       0
                  300    0.000   0.000   0.000   0.000   0.000       0
-----
Flags: (*) % of frames beyond alarm threshold in the current measurement window
Total Configured Segments      : 3
Total Active Segments          : 3

```

History

This command was first available in ExtremeXOS 12.3.

The **mep id** show output was added in ExtremeXOS 15.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show cfm segment frame-delay

```
show cfm segment frame-delay {segment_name}}
```

Description

This command displays frame-delay information for the given CFM segment.

Syntax Description

<i>segment_name</i>	An alpha numeric string identifying the segment name.
---------------------	---

Default

N/A.

Usage Guidelines

Use this command to display frame-delay information for the given CFM segment.

Example

Example output not yet available and will be provided in a future release.

History

This command was first available in ExtremeXOS 12.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show cfm segment frame-delay/frame-loss mep id

```
show cfm segment {{segment_name} | {frame-delay {segment_name}} |  
                  {frame-loss {segment_name {mep mep_id}}}}
```

Description

This command is used to display the current status and configured values of a CFM segment.

Syntax Description

<i>segment_name</i>	An alphanumeric string identifying the segment name.
---------------------	--

Default

N/A.

Usage Guidelines

Use this command to display the current status and configured values of a CFM segment.



Note
In this command, the row “pending frames” will be displayed only for on-demand mode of transmission.

A segment is considered as active if any of the MEPs in the segment is enabled for Frame Loss measurement. Active Segment count will be incremented by one only even if there are multiple MEPs enabled for Frame Loss. For example, assume that there are 3 segments created - seg1, seg2 and seg3.

Segment "seg1" is enabled for Frame Delay measurement. Segment "seg3" has 10 MEPs added with 4 enabled for Frame Loss measurement, the following are the valid counts. Switch wide "Total Configured Segments" will be 3 and "Total Active Segments" will be 2. For Segments "seg1" and "seg2", "Total Configured MEPs" and "Total Active MEPs" will be 0. For segment "seg3", "Total Configured MEPs" will be 10 and "Total Active MEPs" will be 4.

By default, both the Frame Delay and Frame Loss sections are displayed for all the CFM segments. The user has option to filter out based on Segment Name or Frame Delay / Frame Loss.

The behavior for each of the optional parameters is explained below:

- `show cfm segment`: Displays frame-delay and frame-loss information for all the CFM segments.
- `show cfm segment segment_name`: Displays frame-delay and frame-loss information for the given CFM segment.
- `show cfm segment frame-delay`: Displays frame-delay information for all the CFM segments.
- `show cfm segment frame-delay segment_name`: Displays frame-delay information for the given CFM segment.
- `show cfm segment frame-loss`: Displays frame-loss information for all the CFM segments (and all the MEPs under each of the segment).
- `show cfm segment frame-loss segment_name`: Displays frame-loss information for the given CFM segment (and all the MEPs under the given segment).
- `show cfm segment frame-loss segment_name mep mep_id`: Displays frame-loss information for the given CFM segment - MEP ID combination.

Example

```
Switch#show cfm segment sc-rtp
CFM Segment Name      : sc-rtp
Domain Name           : pbt-d2
Association            : pbt-d2-protecting
MD Level              : 2
Destination MAC        : 00:04:96:1e:14:70
Frame Delay:
MEP ID                : 100
DMM Transmission      : In Progress
Transmission mode     : Continuous
Frames Transmitted    : 24
Frames Received       : 15
DMM Tx Interval       : 2 secs
DMR Rx Timeout        : 10 msec
Alarm Threshold       : 10 %
Clear Threshold       : 95 %
Measurement Window Size : 60
Class of Service      : 0
Tx Start Time         : Fri Apr 17 01:29:45 2009
Min Delay             : Fri Apr 17 01:30:29 2009
Max Delay             : Fri Apr 17 01:30:03 2009
Last Alarm Time       : Fri Apr 17 01:29:59 2009
Alarm State           : Set
Lost Frames in Current Window : 9

MEP ID                : 200
DMM Transmission      : In Progress
Transmission mode     : Continuous
```



```

Frames Transmitted      : 24
Frames Received         : 15
DMM Tx Interval         : 2 secs
DMR Rx Timeout          : 10 msec
Alarm Threshold         : 10 %
Clear Threshold         : 95 %
Measurement Window Size : 60
Class of Service        : 0
Tx Start Time           : Fri Apr 17 01:29:45 2009
Min Delay               : Fri Apr 17 01:30:29 2009
Max Delay               : Fri Apr 17 01:30:03 2009
Last Alarm Time         : Fri Apr 17 01:29:59 2009
Alarm State             : Set
Lost Frames in Current Window : 9

Frame Loss:
LMM Tx Interval         : 2 secs
LMR Rx Timeout          : 10 msec
SES Threshold           : 30 %
Consecutive Available Count : 10
Measurement Window Size : 60
Class of Service        : 0
Total Configured MEPs   : 2
Total Active MEPs       : 2
MEP ID                  : 100
LMM Transmission        : In Progress
Transmission mode       : Continuous
Frames Transmitted      : 24
Frames Received         : 15
Availability Status      : Available/Unavailable
Unavailability Start Time : Fri Apr 17 01:10:45 2011
Unavailability End Time  : Fri Apr 17 01:20:45 2011
Tx Start Time           : Fri Apr 17 01:10:45 2011
Min Near-End Frame Loss : Fri Apr 17 01:29:45 2009
Max Near-End Frame Loss : Fri Apr 17 01:39:45 2009
Min Far-End Frame Loss  : Fri Apr 17 01:49:45 2009
Max Far-End Frame Loss  : Fri Apr 17 01:59:45 2009
MEP ID                  : 200
LMM Transmission        : In Progress
Transmission mode       : Continuous
Frames Transmitted      : 24
Frames Received         : 15
Availability Status      : Available/Unavailable
Unavailability Start Time : Fri Apr 17 01:10:45 2011
Unavailability End Time  : Fri Apr 17 01:20:45 2011
Tx Start Time           : Fri Apr 17 01:10:45 2011
Min Near-End Frame Loss : Fri Apr 17 01:29:45 2009
Max Near-End Frame Loss : Fri Apr 17 01:39:45 2009
Min Far-End Frame Loss  : Fri Apr 17 01:49:45 2009
Max Far-End Frame Loss  : Fri Apr 17 01:59:45 2009
-----
Total Configured Segments : 1
Total Active Segments    : 1

```

History

This command was first available in ExtremeXOS 12.3.

The **mep id** show output was added in ExtremeXOS 15.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show cfm segment frame-loss statistics

```
show cfm segment frame-loss statistics {segment-name}
```

Description

Displays shows frame-loss statistics.

Syntax Description

<i>segment_name</i>	An alpha numeric string identifying the segment name.
---------------------	---

Default

N/A.

Usage Guidelines

The below output is an example for displaying the frame-loss stats for the cfm segments. This command shows the recent, minimum, maximum and average near-end and far-end frame loss ratios during the current transmission. The stats for a particular segment will be preserved till the user triggers the next LMM transmission or until it does a clear counter.

Example

The following command displays the frame loss statistics for the CFM segment:

```
LEFT.93 # show cfm segment frame-loss statistics
-----
Segment Name      MEP      Last      Last      Min      Max      Min      Max      Mean
Mean
                  ID       NE       FE       NE       NE       FE       FE       NE
FE
                  FLR      FLRFLRFLRFLRFLRFLRFLR      NLR
-----
seg1              111      10      10      10      10      10      10      10      10
seg1              222      10      10      10      10      10      10      10      10
seg2              333      10      10      10      10      10      10      10      10
-----
Legend: FE - Far End, NE - Near End, FLR - Frame Loss Ratio

Window FE FLR   Last FE Tx   Last FE Rx
-----
cs2              3      0.000000e+00   509467221   526672689
0.000000e+00   501936465   544907407
```

 Legend: FE - Far End, NE - Near End, FLR - Frame Loss Ratio

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show cfm segment frame-loss

```
show cfm segment frame-loss {segment_name}
```

Description

This command displays frame-loss information for the given CFM segment.

Syntax Description

<i>segment_name</i>	An alpha numeric string identifying the segment name.
---------------------	---

Default

N/A.

Usage Guidelines

Use this command to display frame-delay information for the given CFM segment.

Example

```
sho cfm seg frame-loss
CFM Segment Name      : cs2
Domain Name           : dom2
Association            : a2
MD Level              : 2
Destination MAC        : 00:04:96:52:a7:64
Frame Loss:
LMM Tx Interval       : 10 secs
SES Threshold         : 1.000000e-02
Consecutive Available Count : 4
Measurement Window Size : 1200
Class of Service      : 6
Total Configured MEPs : 1
Total Active MEPs     : 1
MEP ID                : 3
LMM Transmission      : In Progress
```

```

Transmission Mode      : Continuous
Frames Transmitted     : 483
Frames Received        : 483
Availability Status    : Available
Unavailability Start Time : None
Unavailability End Time : None
Tx Start Time          : Mon Apr 23 12:28:28 2012
-----
Total Configured Segments : 1
Total Active Segments    : 1
#
#

```

History

This command was first available in ExtremeXOS 12.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show cfm segment mep

```
show cfm segment {segment_name} {mep mep_id }
```

Description

This command displays frame-delay information for the given CFM segment – MEP ID combination.

Syntax Description

<i>segment_name</i>	An alpha numeric string identifying the segment name.
---------------------	---

Default

N/A.

Usage Guidelines

Use this command to display frame-delay information for the given CFM segment – MEP ID combination.

Example

```

Switch#showcfm segment sc-rtp
CFM Segment Name      : sc-rtp
Domain Name           : pbt-d2
Association            : pbt-d2-protectingMD
Level                 : 2
Destination MAC        : 00:04:96:1e:14:70
Frame Delay:

```

```

MEP ID                               : 100

DMM Transmission                     : In Progress
  Transmission mode                   : Continuous
  Frames Transmitted                   : 24
  Frames Received                     : 15
  DMM Tx Interval                     : 2 secs
  DMR Rx Timeout                      : 10 msec
  Alarm Threshold                     : 10 %
  Clear Threshold                     : 95 %
  Measurement Window Size             : 60
  Class of Service                    : 0
  Tx Start Time                      : Fri Apr 17 01:29:45 2009
  Min Delay                          : Fri Apr 17 01:30:29 2009
  Max Delay                          : Fri Apr 17 01:30:03 2009
  Last Alarm Time                    : Fri Apr 17 01:29:59 2009
  Alarm State                        : Set
  Lost Frames in Current Window      : 9

MEP ID                               : 200
DMM Transmission                     : In Progress
  Transmission mode                   : Continuous
  Frames Transmitted                   : 24
  Frames Received                     : 15
  DMM Tx Interval                     : 2 secs
  DMR Rx Timeout                      : 10 msec
  Alarm Threshold                     : 10 %
  Clear Threshold                     : 95 %
  Measurement Window Size             : 60
  Class of Service                    : 0
  Tx Start Time                      : Fri Apr 17 01:29:45 2009
  Min Delay                          : Fri Apr 17 01:30:29 2009
  Max Delay                          : Fri Apr 17 01:30:03 2009
  Last Alarm Time                    : Fri Apr 17 01:29:59 2009
  Alarm State                        : Set
  Lost Frames in Current Window      : 9

Frame Loss:
  LMM Tx Interval                    : 2 secs
  LMR Rx Timeout                     : 10 msec
  SES Threshold                      : 30 %
  Consecutive Available Count        : 10
  Measurement Window Size            : 60
  Class of Service                   : 0
  Total Configured MEPs              : 2
  Total Active MEPs                  : 2

MEP ID                               : 100
LMM Transmission                     : In Progress
  Transmission mode                   : Continuous
  Frames Transmitted                   : 24
  Frames Received                     : 15
  Availability Status                 : Available/Unavailable
  Unavailability Start Time           : Fri Apr 17 01:10:45 2011
  Unavailability End Time             : Fri Apr 17 01:20:45 2011      Tx
    Start Time                       : Fri Apr 17 01:10:45 2011      Min
    Near-End Frame Loss               : Fri Apr 17 01:29:45 2009      Max
    Near-End Frame Loss               : Fri Apr 17 01:39:45 2009      Min
    Far-End Frame Loss                : Fri Apr 17 01:49:45 2009      Max
    Far-End Frame Loss                : Fri Apr 17 01:59:45 2009

MEP ID                               : 200
LMM Transmission                     : In Progress
  Transmission mode                   : Continuous
  Frames Transmitted                   : 24

```

```
Frames Received           : 15
Availability Status       : Available/Unavailable
Unavailability Start Time : Fri Apr 17 01:10:45 2011
Unavailability End Time   : Fri Apr 17 01:20:45 2011
Tx Start Time            : Fri Apr 17 01:10:45 2011
Min Near-End Frame Loss   : Fri Apr 17 01:29:45 2009
Max Near-End Frame Loss   : Fri Apr 17 01:39:45 2009
Min Far-End Frame Loss    : Fri Apr 17 01:49:45 2009
Max Far-End Frame Loss    : Fri Apr 17 01:59:45 2009

-----
Total Configured Segments : 1
Total Active Segments    : 1
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show cfm segment

```
show cfm segment {segment_name}
```

Description

Displays information for CFM segments.

Syntax Description

<i>segment_name</i>	An alpha numeric string identifying the segment name.
---------------------	---

Default

N/A.

Usage Guidelines

Use this command to display information for the selected CFM segment.

If a segment name is not specified, the information for all of the segments that are currently configured are displayed.

Example

The following command displays information for an active CFM segment that is configured to transmit with a specific count:

```
show cfm segment s2
CFM Segment Name      : s2
Domain Name           : pbt-d2
Association            : pbt-d2-protecting
MD Level              : 2
Destination MAC        : 00:04:96:1e:14:70
DMM Transmission      : In Progress
Transmission mode      : Continuous
Frames Transmitted     : 2
Frames Received        : 2
DMM TX Interval       : 2secs
DMR RX Timeout        : 10 msec
Alarm Threshold        : 10 %
Clear Threshold        : 95 %
Measurement Window Size : 60
Class of Service       : 0
Tx Start Time          : Sun Apr 19 21:18:58 2009
Min Delay              : Sun Apr 19 21:18:58 2009
Max Delay              : Sun Apr 19 21:19:00 2009
Last Alarm Time        : None
Alarm State            : Not Set
Lost Frames in Current window : 0
-----
Total Configured Segments : 2
Total Active Segments    : 1
```

The following command displays information for a disabled segment:

```
BD-12804.1 # sh cfm seg s2
CFM Segment Name      : s2
Domain Name           : pbt-d2
Association            : pbt-d2-protecting
MD Level              : 2
Destination MAC        : 00:04:96:1e:14:70
DMM Transmission      : Disabled
Frames Transmitted     : 10
Frames Received        : 10
DMM TX Interval       : 2secs
DMR RX Timeout        : 10 msec
Alarm Threshold        : 10 %
Clear Threshold        : 95 %
Measurement Window Size : 60
Class of Service       : 0
Tx Start Time          : Sat Apr 18 05:39:54 2000
Min Delay              : Sat Apr 18 05:40:12 2000
Max Delay              : Sat Apr 18 05:39:56 2000
Last Alarm Time        : None
Alarm State            : Not Set
Lost Frames in Current window : 1
-----
Total Configured Segments : 2
Total Active Segments    : 0
```

History

This command was first available in ExtremeXOS 12.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show cfm session counters missed-hellos

```
show cfm session counters missed-hellos { domain_name { association_name
    { {ports port_list} { end-point [up|down] } } } } {history | no-
    refresh | refresh}
```

Description

This command displays current and historical CFM session missed-hellos statistics.

Syntax Description

<i>domain_name</i>	IEEE 802.1ag domain name.
<i>association_name</i>	IEEE 802.1ag or ITU-T Y.1731 association name.
ports	Specify ports to show.
<i>port_list</i>	List of ports to show.
end-point	Show MEPs (Maintenance association End Point).
up	End point is up.
down	End point is down.
history	Historical set of bins.
no-refresh	Page by page display without continuous refresh.
refresh	Continuous refresh of output.

Default

Refresh.

Usage Guidelines

None.

Example

The following example show the current statistics with no-refresh:

```
# show cfm session counters missed-hellos no-refresh
      1      2      3      4      5      6      7      8
123456789012345678901234567890123456789012345678901234567890
=====
```



```

Session ID      Port  Remote End-Point  Interval  Missed  Missed  Down  Flags
                MAC Address      (msec)    Once    Twice
=====
3-4000-1111-6666 4      00:11:22:33:44:55    3.3      ---     ---    >9999  DHIs
3-4000-2222-6666 6      00:11:22:33:44:55   10000    >9999    >9999    >9999  DSIs
3-4001-3333-7777 5      00:66:77:88:99:aa   600000    >9999    >9999    >9999  USAs
=====
Session ID: MD Level-VLAN ID-Local MEP ID-Remote MEP ID
Flags: Maintenance Point: (U) Up, End-Point, (D) Down End-Point
      Session Type: (S) Software, (H) Hardware
      Status: (A) Active, (I) Inactive
      Remote End-Point MAC Address: (d) dynamic, (s) static

```

The following example displays current statistics with refresh.

```

# show cfm session counters missed-hellos
=====
Session ID      Port  Remote End-Point  Interval  Missed  Missed  Down  Flags
                MAC Address      (msec)    once    twice
=====
3-4000-1111-6666 4      00:11:22:33:44:55    3.3      ---     ---    >9999  DHIs
3-4000-2222-6666 6      00:11:22:33:44:55   10000    >9999    >9999    >9999  DSIs
3-4001-3333-7777 5      00:66:77:88:99:aa   600000    >9999    >9999    >9999  USAs
=====
Session ID: MD Level-VLAN ID-Local MEP ID-Remote MEP ID
Flags: Maintenance Point: (U) Up End-Point, (D) Down End-Point
      Session Type: (S) Software, (H) Hardware
      Status: (A) Active, (I) Inactive
      Remote End-Point MAC Address: (d) dynamic, (s) static
      0->Clear Counters  U->page up  D->page down  ESC->exit
#

```

History

This command was first available in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show cfm

```

show cfm { domain_name { association_name {{ports port_list
      {[intermediate-point | [end-point [up|down]]}}}

```

Description

Displays the current CFM configuration on the switch.

Syntax Description

<i>domain_name</i>	Enter the name of the domain you want to display.
<i>association_name</i>	IEEE 802.1ag or ITU-T Y.1731 association name.
<i>port_list</i>	Enter the ports in the domain and association you want to display.
up	Enter this to display the UP MEP for the specified MA.
down	Enter this to display the DOWN MEP for the specified MA.
intermediate-point	Enter this to display the MIPs for the specified MA.

Default

N/A.

Usage Guidelines

This command displays the following information:

- Domain names
- MA levels
- Association names
- VLAN names
- Transmit Interval
- UP MEPs
- MEPIDs
- MEP transmit intervals
- MEP State
- DOWN MEPs
- Intermediate points (MIPs)
- Total number of CFM ports on the switch
- Destination MAC Type
- VPLS-based MPs
- Sender ID information
- ISID Intermediate Point

For the number of domains, ports, MEPs, MIPs, and associations supported on the switch, see the *Supported Instances for CFM* section in the [Switch Engine 32.3 User Guide](#).

Example

The `show cfm` command displays the current CFM configuration on the switch:

```
* switch # show cfm
Domain: "dnsname", MD Level: 2
  Association: "10", Destination MAC Type: Multicast, VLAN "v1" with 2 cfm ports
  Transmit Interval: 1000 ms, Type : IEEE 802.1ag Maintenance Association
    port 1; Up End Point, mepid: 100, transmit-interval: 1000 ms (from association)
```

```

MEP State: Enabled, CCM Message: Enabled, Send SenderId TLV:
Disabled
    Faulting State                : No
    Last Faulting State Change    : Wed Jun 19 09:12:13 2013
    MEP Error Defects             : None
    Port Status                   : Up
    port 15; Intermediate Point ( Dynamic )
    Association: "VSNLMEG1", Destination MAC Type: Multicast, VLAN "none" with 0 cfm
ports
    Transmit Interval: 1000 ms, Type : ITU-T Y.1731 Maintenance Entity Group
    Association: "snmp_ma_name", Destination MAC Type: Multicast, VLAN "none" with 0
cfm ports
    Transmit Interval: 1000 ms, Type : IEEE 802.1ag Maintenance Association
    Domain: "dom1", MD Level: 5
    Association: "VSNLMEG1", Destination MAC Type: Multicast, VLAN "v2" with 2 cfm
ports
    Transmit Interval: 10 ms, Type : ITU-T Y.1731 Maintenance Entity Group
    port 1; Down End Point, mepid: 100, transmit-interval: 3.3 ms
(configured)
MEP State: Enabled, CCM Message: Enabled, Send SenderId TLV:
Disabled
    Faulting State                : Yes
    Last Faulting State Change    : Wed Jun 19 09:08:12 2013
    MEP Error Defects             : Remote
    Port Status                   : Up
    port 15; Down End Point, mepid: 200, transmit-interval: 10 ms (from association)
MEP State: Enabled, CCM Message: Enabled, Send SenderId TLV:
Disabled
    Faulting State                : Yes
    Last Faulting State Change    : Wed Jun 19 09:08:13 2013
    MEP Error Defects             : Remote
    Port Status                   : Up
    Association: "short_ma_name", Destination MAC Type: Multicast, VLAN "v1" with 2
cfm ports
    Transmit Interval: 1000 ms, Type : IEEE 802.1ag Maintenance Association
    port 1; Down End Point, mepid: 100, transmit-interval: 3.3 ms
(configured)
MEP State: Enabled, CCM Message: Enabled, Send SenderId TLV:
Disabled
    Faulting State                : Yes
    Last Faulting State Change    : Wed Jun 19 09:09:47 2013
    MEP Error Defects             : Remote
    Port Status                   : Up
    port 15; Down End Point, mepid: 200, transmit-interval: 1000 ms (from
association)
MEP State: Enabled, CCM Message: Enabled, Send SenderId TLV:
Disabled
    Faulting State                : Yes
    Last Faulting State Change    : Wed Jun 19 09:09:47 2013
    MEP Error Defects             : RDI, Remote
    Port Status                   : Up
    Domain: "dom2", MD Level: 6
    Association: "VSNLMEG1", Destination MAC Type: Unicast, VLAN "v2" with 2 cfm
ports
    Transmit Interval: 1000 ms, Type : ITU-T Y.1731 Maintenance Entity Group
    port 1; Up End Point, mepid: 100, transmit-interval: 1000 ms (from
association)
MEP State: Enabled, CCM Message: Enabled, Send SenderId TLV:
Disabled
    Faulting State                : No
    Last Faulting State Change    : Wed Jun 19 09:39:14 2013
    MEP Error Defects             : None
    Port Status                   : Up
    port 15; Intermediate Point ( Dynamic )Domain: "00:11:22:33:44:55.6666", MD

```

```

Level: 7
    Association: "VSNLMEG1", Destination MAC Type: Multicast, VLAN "v3" with 2 cfm
ports
    Transmit Interval: 1000 ms, Type : ITU-T Y.1731 Maintenance Entity Group
port 1; Up End Point, mepid: 100, transmit-interval: 1000 ms (from
association)
    MEP State: Enabled, CCM Message: Enabled, Send SenderId TLV:
Disabled
    Faulting State : No
    Last Faulting State Change : Wed Jun 19 09:10:10 2013
    MEP Error Defects : None
    Port Status : Up
port 15; Intermediate Point ( Dynamic )
    Association: "short_ma_name", Destination MAC Type: Multicast, VLAN "v4" with 2
cfm ports
    Transmit Interval: 1000 ms, Type : IEEE 802.1ag Maintenance Association
port 1; Up End Point, mepid: 100, transmit-interval: 1000 ms (from
association)
    MEP State: Enabled, CCM Message: Enabled, Send SenderId TLV:
    Enabled IPaddress:
10.10.10.1
    Faulting State : No
    Last Faulting State Change : Wed Jun 19 09:15:08 2013
    MEP Error Defects : None
    Port Status : Up
port 15; Intermediate Point ( Dynamic )
Total Number of Domain : 4
Total Number of Association : 8
Total Number of Up MEP : 4
Total Number of Down MEP : 4
Total Number of MIP : 4
Total Number of CFM port : 12
Total Number of SW MEP : 4
Total Number of HW MEP : 4
Total Number of VPLS MIP(Static/Up): 0 / 0

=====
MEP Error Defect Types:
    Remote : Not receiving CCMs from Remote MEP
    Error : Erroneous CCM received
    XCON : Cross-connect CCM received
    RDI : Remote Defect Indication sent by some MEP

```

History

This command was first available in ExtremeXOS 11.4.

Transmit Interval and MEP State were added in ExtremeXOS 12.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show checkpoint-data

```
show checkpoint-data {process}
```

Description

Displays the status of one or more processes being copied from the master node to the backup node.

Syntax Description

<i>process</i>	Specifies the name of the processes being copied.
----------------	---

Default

N/A.

Usage Guidelines

This command displays, in percentages, the amount of internal state copying completed by each process and the traffic statistics between the process on both the master and the backup nodes.

This command is also helpful in debugging synchronization problems that occur at run-time. To check the status of synchronizing the nodes, use the [show switch](#) command.

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

Example

The following command displays the checkpointing status and the traffic statics of all of the processes between the master and the backup nodes:

```
# show checkpoint-data
Process      Tx      Rx  Errors   Sent   Total    % Chkpt   Debug-info
-----
devmgr       3812   1731      0      3      3   100% ON OK 1 (00008853)
dirser        0      0      0      0      0    0% ON OK 1 (000008D3)
ems           5      0      0      0      0   100% ON OK 1 (000008D3)
nodemgr       0      0      0      0      0    0% ON OK 1 (000008D3)
snmpSubagent  0      0      0      0      0    0% ON OK 1 (000018D3)
snmpMaster    0      0      0      0      0    0% ON OK 1 (000008D3)
cli           0      0      0      0      0    0% ON OK 1 (000018D3)
edp           0      0      0      0      0    0% ON OK 1 (000008D3)
cfgmgr       82      82      0      1      1   100% ON OK 1 (000018D3)
elrp          0      0      0      0      0    0% ON OK 1 (000008D3)
vlan        1047      1      0      0      0   100% ON OK 1 (000008D3)
aaa           0      0      0      0      0    0% ON OK 1 (000008D3)
fdb          957      2      0      0      0   100% ON OK 1 (000008D3)
msgsrv        0      0      0      0      0   100% ON OK 1 (000008D3)
eaps          0      0      0      0      0    0% ON OK 1 (000008D3)
stp           1      0      0      0      0    0% ON OK 1 (000008D3)
esrp          1      0      0      0      0   100% ON OK 1 (000008D3)
polMgr        0      0      0      0      0    0% ON OK 1 (000008D3)
mcmgr         2      2      0      0      0   100% ON OK 1 (000008D3)
acl           0      0      0      0      0   100% ON OK 1 (000008D3)
netLogin      0      0      0      0      0    0% ON OK 1 (000008D3)
ospf          0      0      0      0      0    0% ON OK 1 (000008D3)
netTools      1      0      0      0      0   100% ON OK 1 (000008D3)
telnetd       0      0      0      0      0    0% ON OK 1 (000008D3)
rtmgr         4      4      0      0      0   100% ON OK 1 (000008D3)
```

```

vrrp          378      0      0      0      0      0% ON OK 1 (000008D3)
tftpd         0        0      0      0      0      0% ON OK 1 (000008D3)
thttpd        0        0      0      0      0      0% ON OK 1 (000008D3)
rip           0        0      0      0      0      0% ON OK 1 (000008D3)
dosprotect    0        0      0      0      0      0% ON OK 1 (000008D3)
epm           0        0      0      0      0      0% ON OK 1 (000008D3)
hal           0        0      0      0      0      0% ON OK 1 (000008D3)
bgp           0        0      0      0      0      0% ON OK 1 (000008D3)
pim           0        0      0      0      0      0% ON OK 1 (000008D3)
etmon         185     185      0      0      0 100% ON OK 1 (000008D3)
Flags : S - Server started, c - Client started, D - Checkpoint dependency satisfied
        C - Checkpointing is ON, f - No config dependency, N - DMLIB not in sync
        R - CM Backend not ready, l - Process is loading config, L - Process config is
loaded
        I - IPML connection alive, n - Backup process not in sync with Primary process

```

To view the output for a specific process, use the process option. The following command displays detailed information for the *STP* process:

```

# show checkpoint-data stp
Process      Tx      Rx  Errors   Sent   Total      % Chkpt   Debug-info
-----
stp          1        0      0      0      0      0% ON OK 1 (000008D3)

```

History

This command was first available in ExtremeXOS 10.1.

An error count was added to the output in ExtremeXOS 11.1.

IPML connection alive flag added in ExtremeXOS 22.5.

The "n" flag for the backup process being out of sync was added in ExtremeXOS 31.1.

Platform Availability

This command is available only on SummitStack.

show clear-flow

```
show clear-flow
```

Description

Displays the status of the CLEAR-Flow agent, any CLEAR-Flow policies on each interface, and the number of CLEAR-Flow rules.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following display shows output for the command show clear-flow:

```
clear-flow: Enabled
VLAN      Port  Policy Name      No. of CF Rules
=====
*         2:1   CFexample        6
*         2:26  CFexample        6
*         2:40  CFexample        6
Default   *     CFexample        6
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

VLAN

show clear-flow acl-modified

```
show clear-flow acl-modified
```

Description

Displays the ACLs modified by CLEAR-Flow actions.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command displays the ACLs that have been modified by CLEAR-Flow rules that have been triggered.

Example

The following display shows output for the command `show clear-flow acl-modified`:

```

Policy Name      Vlan Name      Port Rule Name      Default ACL      CF Added
Actions          Actions
=====
clearFlow        *              2:26 acl-rule-4      D                QP1
clearFlow        *              2:26 acl-rule-3      D                D
clearFlow        *              2:26 acl-rule-2      D                M
clearFlow        *              2:26 acl-rule-1      D                P
clearFlow        Default        *      acl-rule-4      D                QP1
clearFlow        Default        *      acl-rule-3      D                D
clearFlow        Default        *      acl-rule-2      D                M
clearFlow        Default        *      acl-rule-1      D                P
=====
Total Entries: 8
Notation:
P - Permit, D- Deny, M - mirror enabled, m - mirror disabled

```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

VLAN

show clear-flow rule

```

show clear-flow [port port | vlan vlanname | any] [rule rulename]
                {detail}

```

Description

Displays the CLEAR-Flow rules, values, and configuration.

Syntax Description

<i>port</i>	Specifies the port.
<i>vlanname</i>	Specifies the <u>VLAN</u> .
any	Specifies the wildcard interface.
<i>rulename</i>	Specifies the entry name of a CLEAR-Flow rule.
detail	Display detailed information.

Default

N/A.

Usage Guidelines

If you issue the command without the rule keyword, all of the CLEAR-Flow rules for the policy on the port, VLAN, and the wildcard are displayed. If you specify a rule name, only that rule will be displayed. The detail keyword displays detailed information about the rule.

Example

The following display shows output for the command `show clear-flow port 2:6`:

Rule Name Value	Type Oper	Period	Last	Rel If	Threshold Else	TCNT	NumAction
rule-count	CN	30	16892762	>	100	7	3 3
rule-delta	DT	30	7762385	>	1000	1	4 3
rule-delta-2	DT	5	0	>	1000	0	4 3
rule-delta-ratio	DR	30	0	>	20	0	2 0
rule-ratio	RT	30	0	>	10	0	3 3
rule-ratio-2	RT	5	0	>	10	0	3 3
=====							
Total Entries: 6							
Notation:							
Threshold Type: CN - Count, DT - Delta, RT - Ratio, DR - DeltaRatio							
TCNT - Number of times expression is continuously evaluated to be true							

The following display shows output for the command `show clear-flow port 2:6 rule rule-delta detail`:

```

Rule Name: rule-delta      Sample Period: 30      Hysteresis: 20
=====
DELTA(counter1) = 0 sampled at 24 seconds ago
Expression evaluation is currently FALSE
if (DELTA(counter1) > 1000) then {
PERMIT:    Allow ACL rule acl-rule-3
SYSLOG:    [INFO] [Delta $ruleValue counter $counter1 offset $counterOffset1 delTime
$deltaTime delay $delayTime]
CLI:       [disable port $port]
QOS:       Set rule acl-rule-4 qos value to QP6
} else {
DENY:     Block ACL rule acl-rule-3
QOS:       Set rule acl-rule-4 qos value to QP1
CLI:       [enable port $port]
}

```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

VLAN

show clear-flow rule-all

```
show clear-flow rule-all
```

Description

Displays all the CLEAR-Flow rules on the switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following display shows output for the command `show clear-flow rule-all`:

```

Policy Name      Vlan Name      Port Rule Name  Last Value OP Threshold  TCNT  Sec
=====
clearFlow        *              2:1  rule-count  1          > 100      0    11
clearFlow        *              2:1  rule-delta  1          > 1000    0    11
clearFlow        *              2:1  rule-delta  0          > 1000    0     4
clearFlow        *              2:1  rule-delta  0          > 20      0    11
clearFlow        *              2:1  rule-ratio  0          > 10      0    11
clearFlow        *              2:1  rule-ratio  0          > 10      0     4
clearFlow        *              2:26 rule-count  9030635    > 100     1    10
clearFlow        *              2:26 rule-delta  9030635    > 1000    1    10
clearFlow        *              2:26 rule-delta  0          > 1000    0     4
clearFlow        *              2:26 rule-delta  0          > 20      0    10
clearFlow        *              2:26 rule-ratio  0          > 10      0    10
clearFlow        *              2:26 rule-ratio  0          > 10      0     4
clearFlow        Default        *      rule-count  36666439   > 100     1    10
clearFlow        Default        *      rule-delta  36666439   > 1000    1    10
clearFlow        Default        *      rule-delta  0          > 1000    0     4
clearFlow        Default        *      rule-delta  0          > 20      0    10
clearFlow        Default        *      rule-ratio  0          > 10      0    10
clearFlow        Default        *      rule-ratio  0          > 10      0     4
=====
Total Entries: 18
Notation:
TCNT - Number of times expression is continously evaluated to be true
Sec - Number of seconds elapsed from last sampled data

```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

VLAN

show clear-flow rule-triggered

```
show clear-flow rule-triggered
```

Description

Displays the triggered CLEAR-Flow rules.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command displays the rules that have been triggered; in other words, the rule threshold has been reached.

Example

The following display shows output for the command show clear-flow rule-triggered:

```
Policy Name      Vlan Name      Port Rule Name  Last Value OP Threshold  TCNT  Sec
=====
clearFlow        *              2:26 rule-count  9130377   >   100        2    25
clearFlow        *              2:26 rule-delta  99742     >  1000        2    25
clearFlow        Default        *      rule-count  37069465  >   100        2    25
clearFlow        Default        *      rule-delta  403026    >  1000        2    25
=====
Total Entries: 4
Notation:
TCNT - Number of times expression is continuously evaluated to be true
Sec - Number of seconds elapsed from last sampled data
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

VLAN

show cli journal

```
show cli journal
```

Description

This command shows the historical list (journal) of the most recently executed CLI commands.

Syntax Description

This command has no arguments or variables.

Default

By default, one hundred commands are preserved in the journal.

Usage Guidelines

The journal retains as many as 200 of the most recently executed commands along with the timestamp and user name. Commands are saved even after logging off, rebooting, or switch crashes.

To set the size (number of commands) of the journal, use the configure `cli journal size size` command.

Example

The following is sample output from the command:

Timestamp	Session	User Name	Command

03/31/2016 13:29:29.62	serial	admin	show tech
03/31/2016 13:29:30.21	serial	unknown	debug cfgmgr show next
maximum-rows 1 vlan.show_ports_info port=None portList=*			
03/31/2016 13:29:30.28	serial	unknown	debug cfgmgr show next
poe.poe_extremePethPseSlotEntry			
03/31/2016 13:29:30.40	serial	unknown	show switch
03/31/2016 13:29:30.50	serial	unknown	show version detail
03/31/2016 13:29:30.57	serial	unknown	show version images
03/31/2016 13:29:30.64	serial	unknown	show management
03/31/2016 13:29:30.78	serial	unknown	show session
03/31/2016 13:29:30.85	serial	unknown	show license
03/31/2016 13:29:30.92	serial	unknown	ls
03/31/2016 13:29:31.00	serial	unknown	ls internal-memory
03/31/2016 13:29:31.08	serial	unknown	debug hal show compact-
flash			
03/31/2016 13:29:31.15	serial	unknown	show odometers
03/31/2016 13:29:31.23	serial	unknown	show fans detail
03/31/2016 13:29:31.29	serial	unknown	show temperature
03/31/2016 13:29:31.49	serial	unknown	show power
03/31/2016 13:29:31.56	serial	unknown	show power detail
03/31/2016 13:29:31.62	serial	unknown	show cpu-monitoring
03/31/2016 13:29:33.24	serial	unknown	show memory
03/31/2016 13:29:34.68	serial	unknown	run script mem-
stats.py			
03/31/2016 13:29:35.08	serial	unknown	show edp ports all

History

This command was first available in ExtremeXOS 22.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show configuration

```
show configuration {module-name} {detail}
```

Description

Displays the current configuration for the system or the specified module.

Syntax Description

<i>module-name</i>	Specifies the name of configuration module. The term configuration module refers to feature in ExtremeXOS. By displaying a module, you can view the commands used to configure that feature. For example, to display all of the configurations that you made for only <i>STP</i> , specify the stp as the module-name.
detail	Displays configuration data including default. If the detail option is not specified, only the configuration changes you made to the factory defaults are shown.

Default

N/A.

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.

These files have the .cfg file extension. Do not use a text editor to view or modify your XML-based switch configuration files.

To save the configuration file as an ASCII-formatted file, and to view it with a text editor, see the `upload configuration [hostname | ipaddress] filename {vr vr-name}` and the `load script filename {arg1} {arg2} ... {arg9}` commands.

Beginning with ExtremeXOS 12.1, when you specify `show configuration` only, the switch displays configuration information for each of the switch modules excluding the default data.

You can display only the configuration of a module of interest by using the module-name keyword. For example, some of the modules are AAA, *ACL*, *BGP*, *EDP*, *FDB*, *SNMP*, and *VLAN*. Use [TAB]-completion to see a list.

You must have administrator access to view the output of the `show configuration` command.

Depending on the software version running on your switch, the configurations on your switch, and the type of switch you have, additional or different configuration information may appear.

Loading ARP and neighbor discovery (ND) configurations from versions earlier than ExtremeXOS 30.1 after upgrading to ExtremeXOS 30.1, abide by the following rules:

- The number of static ARP/ND entries configured in all VRs (default + Mgmt + user VRs) do not exceed the new global max. entries.
- The number of static proxy entries configured in all VRs do not exceed the new global max. proxy entries.
- The sum of (per VR configured limits of VR-default + default limit of VR-Mgmt) do not exceed the respective new max. global limit.
- ExtremeXOS ignores pre-30.1 user VR configured limits in 30.1.
- ExtremeXOS logs any static entry restoration failures after they exceed new global limits.

The output of `show configuration fdb` after upgrading to ExtremeXOS 30.1 differs from 22.x output because of the changes in per VR-based IPARP and ND configurations (see examples below).

Example

The following example shows the current configuration of the *OSPF* module on the switch:

```
# show configuration ospf
# Module ospf configuration.
#
configure ospf routerid automatic
configure ospf spf-hold-time 3
configure ospf metric-table 10M 10 100M 5 1G 4 10G 2
configure ospf lsa-batch-interval 30
configure ospf import-policy none
configure ospf ase-limit 0
disable ospf originate-default
disable ospf use-ip-router-alert
disable ospf
configure ospf restart none
configure ospf restart grace-period 120
disable ospf export direct
disable ospf export static
disable ospf export rip
disable ospf export e-bgp
disable ospf export i-bgp
configure ospf area 0.0.0.0 external-filter none
configure ospf area 0.0.0.0 interarea-filter none
```

The following example illustrates the difference between ExtremeXOS 30.1 versus 22.x output for *FDB*.

Example 1

22.x

```
# show configuration fdb detail
configure iparp vr VR-Default max_entries 8192
configure iparp vr VR-Mgmt max_entries 4096
configure neighbor-discovery vr VR-Default max_entries 4096
configure neighbor-discovery vr VR-Mgmt max_entries 4096
```

30.1 and Later

```
# show configuration fdb detail
configure iparp max_entries 12288 ---- > Default values of (VR-Default + VR-Mgmt) = 8192
```

```
+ 4096
configure neighbor-discovery max_entries 8192
```

Example 1

22.x

```
# show configuration fdb
configure iparp vr VR-Default max_entries 10000
configure iparp vr VR-Default max_pending_entries 10
configure iparp vr VR-Default max_proxy_entries 10
configure iparp vr VR-Mgmt max_entries 2000
configure iparp vr VR-Mgmt max_pending_entries 20
configure iparp vr VR-Mgmt max_proxy_entries 20
configure iparp vr vrl max_entries 3000
configure iparp vr vrl max_pending_entries 30
configure iparp vr vrl max_proxy_entries 30
configure neighbor-discovery vr VR-Default max_entries 10000
configure neighbor-discovery vr VR-Default max_pending_entries 100
configure neighbor-discovery vr VR-Mgmt max_entries 1000
configure neighbor-discovery vr VR-Mgmt max_pending_entries 10
configure neighbor-discovery vr vrl max_entries 60
configure neighbor-discovery vr vrl max_pending_entries 70
```

30.1 and Later

```
# show configuration fdb
configure iparp max_entries 14096 ----- max_entries = (Configured VR-Default
max_entries) + (Default VR-Mgmt max_entries) = 10000 + 4096
configure iparp max_pending_entries 266 ----- max_pending_entries = (Configured VR-
Default max_pending_entries) + (Default VR-Mgmt max_pending_entries) = 10 + 256
configure iparp max_proxy_entries 266 ----- max_proxy_entries = (Configured VR-
Default max_proxy_entries) + (Default VR-Mgmt max_proxy_entries) = 10 + 256
configure neighbor-discovery max_entries 14096 ----- max_entries = (Configured VR-
Default max_entries) + (Default VR-Mgmt max_entries) = 10000 + 4096
configure neighbor-discovery max_pending_entries 1124 ----- max_pending_entries =
(Configured VR-Default max_pending_entries) + (Default VR-Mgmt max_pending_entries) =
100 + 1024
```

Example 1

22.x

```
# show configuration fdb
#
# Module fdb configuration.
#
configure iparp vr VR-Default max_entries 10000
configure iparp vr VR-Mgmt max_entries 1000
configure iparp vr vr-user1 max_entries 2000
configure iparp add 20.0.0.100 vr VR-Default 00:0a:0b:0c:0d:0e
configure iparp add 10.68.13.100 vr VR-Mgmt 00:66:77:88:99:aa
configure iparp add 100.0.0.150 vr vr-user 00:11:22:33:44:55
```

30.1 and Later

```
# show configuration fdb
#
# Module fdb configuration.
#
configure iparp max_entries 14096 --- max. limit of one current global table of kernel
configure iparp add 20.0.0.100 vr VR-Default 00:0a:0b:0c:0d:0e ---
configure iparp add 10.68.13.100 vr VR-Mgmt 00:66:77:88:99:aa | --- All static entries of
```

```
all VRs are restored back as long as they do not exceed current max limit.
configure iparp add 100.0.0.150 vr vr-user 00:11:22:33:44:55 --- Once they exceed the
limit, rest of the entries will be discarded and logged as failures.
```

History

This command was first available in ExtremeXOS 11.0.

The detail variable was added in ExtremeXOS 12.1.

The display of blackhole output was first available in ExtremeXOS 12.1.

Impacts to ARP, ND, and FDB output introduced in ExtremeXOS 30.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show configuration difference

```
show configuration difference { from-config-file {to-config-file} }
    {module-name} {detail} {context lines}
```

Description

Displays the difference between switch configurations.

Syntax Description

difference	Compare one configuration file to another.
<i>from-config-file</i>	Configuration file from which to base the difference (the booted configuration if not specified).
<i>to-config-file</i>	Configuration file to which the base configuration is compared (the running configuration if not specified).
<i>module-name</i>	Configuration module name: aaa, acl, bfd, bgp, brm, cfgmgr, cfm, devmgr, dirser, dosprotect, eaps, edp, elrp, elsm, ems, epm, erps, esrp, ethoam, etmon, exsshd, fdb, hal, hclag, idMgr, ipfix, ipSecurity, larp, lldp, mcmgr, mrp, msdp, netlogin, nettools, nodealias, nodemgr, ntp, ospf, ospfv3, otm, pim, policy, polMgr, pwmib, rip, ripng, rtmgr, snmp, stp, techSupport, telnetd, tftpd, throw, tthttpd, twamp, upm, vlan, vmt, vrrp, vsm, xmlc, xmldb
detail	Specifies showing configuration difference including default configuration choices; If not selected, default configuration information does not appear.
context	Adjusts the number of context lines around the difference output. Default is three lines.
<i>lines</i>	Specifies the number of context lines around the difference output. Default is three lines.

Default

If not specified, *from-config-file* is the booted configuration; if not specified, *to-config-file* is the running configuration.

If not specified, the default is three context lines.

Usage Guidelines

Use this command to compare the difference between two configuration files, or configuration files with the running configuration. When using the `save configuration automatic {every minutes {primary | secondary | existing-config | new-config} | never}` command this can be especially helpful for comparing what changed between the running (or "startup") configuration and the most recently autosaved configuration.

Example

The following example compares the factory default configuration (new switch or after `unconfigure switch {all | erase [all | nvram]}` using the **all** keyword:

```
# show configuration difference
--- Factory Default          Tue Feb 27 19:54:10 2018
+++ Running Configuration    Wed Feb 28 12:37:36 2018
@@ -0,0 +1,212 @@
+#
+# Module devmgr configuration.
+#
+configure snmp sysContact "support@extremenetworks.com, +1 888 257 3000"
+#
+# Module vlan configuration.
+#
+create vlan "VLAN_0008"
+configure vlan VLAN_0008 tag 8
.
.
.
+#
+# Module vsm configuration.
+#
```

The following example compares the `primary.cfg` file with the currently running configuration (default) showing default configuration information (**detail** option) and changing the number of context lines to six (as opposed to the default of three):

```
# show configuration difference primary detail context 6
--- primary.cfg              Mon Dec 4 18:08:59 2017
+++ Running Configuration    Mon Dec 4 18:11:45 2017
@@ -608,13 +609,13 @@
disable cli scripting permanent
configure idletimeout 20
enable idletimeout
configure debug core-dumps internal-memory
disable snmp traps configuration save
disable snmp traps configuration change
-save configuration automatic every 2 primary
+save configuration automatic never

#
# Module dosprotect configuration.
```

```
#
disable dos-protect
configure dos-protect interval 1
```

History

This command was first available in ExtremeXOS 22.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show configuration “xmlc”

```
show configuration "xmlc" {detail | non-persistent {detail}}
```

Description

Displays the configuration of an XMLC module.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Displays the configuration of an XMLC module.

When the detail option is chosen, all configuration data including the default is displayed. Otherwise the default would not be displayed.

When the non-persistent option is chosen, UPM non-persistent configuration data is displayed.

Example

The following command displays the xmlc configuration:

```
show configuration "xmlc" detail
```

Following is sample output from this command:

```
Module xmlc configuration.
#
create xml-notification target test url http://10.255.42.73:9080/axis/services/eventPort
configure xml-notification test queue-size 100
disable xml-notification test
```

```
create xml-notification target test2 url http://10.255.42.48:9080/axis/services/eventPort
configure xml-notification test2 queue-size 100
enable xml-notification test2
create xml-notification target epicenter-target url http://10.255.42.48:8080/xos/
webservice
configure xml-notification target epicenter-target user admin encrypted-auth
YWRtaW46ZXBPY2VudGVy
configure xml-notification epicenter-target queue-size 100
enable xml-notification epicenter-target
configure xml-notification target test add idMgr
configure xml-notification target test2 add idMgr
configure xml-notification target epicenter-target add idMgr
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show cos-index

```
show cos-index {cos_index}
```

Description

This command displays a one-to-one mapping of CoS reference to resource.

Syntax Description

<i>cos-index</i>	Class of Service (CoS) index value.
------------------	-------------------------------------

Default

N/A.

Usage Guidelines

None

Example

```
# show cos-index
COS Index  QOS      TOS  TOS  Ingress
           Profile Value Mask  Meter
-----
0    qp0      --   --   --
1    qp1      64 0xfc ingmeter2
2    qp2      --   --  ingmeter15
3    qp3      --   --  ingmeter2
4    qp4      32 0xfc ingmeter12
```

```

5  qp5      --  --  --
6  qp6      --  --  --
7  qp7      --  --  --
32 qp4      --  --  --
255 qp6     --  --  ingmeter3

```

History

This command was first available in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show counters vr

```
show counters vr {vpn-vrf-name}
```

Description

Displays statistics information about VPN VRF operation.

Syntax Description

<i>vpn-vrf-name</i>	Specifies the name of a VPN VRF.
---------------------	----------------------------------

Default

N/A.

Usage Guidelines

This command displays counters that show:

- The total number of IP unicast and multicast routes.
- Route add operation count.
- Route delete operation count.
- Routes dropped count.



Note

The total route count displayed for this command can exceed the total route count displayed by the `show iproute` command because the `show iproute` command displays either unicast or multicast routes, but not both.

This command is supported only on VPN VRFs.

Example

The following example displays the counters for VPN VRF "red":

```
Switch.19 # show counters vr red
Num of Current Routes: 4           Num of Routes Dropped: 0
Num of Route Add:      12          Num of Route Del:      6
Num of Current Routes: 10          Num of Routes Dropped: 0
Num of Route Add:      5           Num of Route Del:      2
```

History

This command was introduced in ExtremeXOS 15.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show cpu-monitoring

show cpu-monitoring {**process** *name*}

Description

Displays the CPU utilization history of one or more processes.

Syntax Description

<i>name</i>	Specifies the name of the process.
-------------	------------------------------------

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.

By default, CPU monitoring is enabled and occurs every 5 seconds. The default CPU threshold value is 90%.

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

Depending on the software version running on your switch or your switch model, additional or different CPU and process information might be displayed.

SummitStack Only

When you issue the command with out any parameters:

- From the stack manager or backup node, the stack displays CPU utilization history for all the processes running on the master node and the backup node in the Active Topology.
- From the stack manager or a standby node, the stack displays CPU utilization history for all the processes running on the master node and the standby node in the Active Topology.

Modular Switches Only

Reading the Output

The show cpu-monitoring command is helpful for understanding the behavior of a process over an extended period of time. The following information appears in a tabular format:

- Process—The name of the process.
- Range of time (5 seconds, 10 seconds, and so forth)—The CPU utilization history of the process or the system. The CPU utilization history goes back only 1 hour.
- Total User/System CPU Usage—The amount of time recorded in seconds that the process spends occupying CPU resources. The values are cumulative meaning that the values are displayed as long as the system is running. You can use this information for debugging purposes to see where the process spends the most amount of time: user context or system context.

Example

The following example displays CPU utilization on the switch:

```
show cpu-monitoring
```

The following is sample output from a switch:

CPU Utilization Statistics - Monitored every 25 seconds											
Process	5	10	30	1	5	30	1	Max	Total		
secs	secs	secs	min	mins	mins	hour	User/System				
util	util	util	util	util	util	util	CPU Usage				
(%)	(%)	(%)	(%)	(%)	(%)	(%)	(secs)				
System	n/a	n/a	0.0	0.9	0.1	0.2	0.5	34.6			
aaa	n/a	n/a	0.0	0.0	0.0	0.0	0.0	1.8	1.72	0.78	
acl	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	0.40	0.24	
bgp	n/a	n/a	0.0	0.0	0.0	0.0	0.0	12.6	11.18	2.21	
cfgmgr	n/a	n/a	0.0	0.0	0.0	0.0	0.8	39.8	4743.92	3575.79	
cli	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	0.59	0.42	
devmgr	n/a	n/a	0.0	0.0	0.0	0.0	0.0	19.5	74.44	24.52	
dirser	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	
dosprotect	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	0.8	0.12	
eaps	n/a	n/a	0.0	0.0	0.0	0.0	0.1	5.5	36.40	15.41	
edp	n/a	n/a	0.0	0.0	0.0	0.0	0.0	11.1	10.92	3.97	
elrp	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	0.49	0.44	
ems	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	1.19	1.29	
epm	n/a	n/a	0.0	0.0	0.0	0.0	0.0	30.7	48.74	32.93	
esrp	n/a	n/a	0.0	0.0	0.0	0.0	0.0	2.7	0.82	0.45	
etmon	n/a	n/a	0.0	0.0	0.0	0.0	0.5	30.5	4865.78	873.87	
...											

The following is sample truncated output from a stack:

```
Slot-1 stackK.3 # sh cpu-monitoring
CPU Utilization Statistics - Monitored every 20 seconds
-----
```

Card	Process	5	10	30	1	5	30	1	Max	Total
secs	secs	secs	min	mins	mins	hour				
util	util	util	util	util	util	util	User/System			
(%)	(%)	(%)	(%)	(%)	(%)	(%)	CPU Usage			
							(secs)			
Slot-1	System	n/a	n/a	0.0	1.6	0.8	0.5	0.5	2.5	
Slot-6	System	n/a	n/a	0.3	0.9	0.7	0.4	0.5	4.6	
Slot-1	aaa	n/a	n/a	0.0	0.0	0.0	0.0	0.0	3.6	1.22
Slot-1	acl	n/a	n/a	0.0	0.0	0.0	0.0	0.0	1.8	1.8
Slot-1	bgp	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Slot-1	brm	n/a	n/a	0.0	0.0	0.0	0.0	0.0	1.8	0.53
Slot-1	cfgmgr	n/a	n/a	0.1	0.0	0.0	0.0	0.0	0.8	3.18
Slot-1	cli	n/a	n/a	0.9	0.8	0.1	0.0	0.3	97.2	13.7
Slot-1	devmgr	n/a	n/a	0.0	0.0	0.0	0.0	0.0	5.0	1.1
Slot-1	dirser	n/a	n/a	0.0	0.0	0.0	0.0	0.0	5.9	0.0
Slot-1	dosprotect	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	0.18
Slot-1	eaps	n/a	n/a	0.0	0.0	0.0	0.0	0.0	1.3	0.92
Slot-1	edp	n/a	n/a	0.0	0.0	0.0	0.0	0.0	1.3	0.68
Slot-1	elrp	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.9	0.49
Slot-1	elsm	n/a	n/a	0.0	0.0	0.0	0.0	0.0	1.3	0.38
Slot-1	ems	n/a	n/a	0.0	0.0	0.0	0.0	0.0	1.3	1.0
Slot-1	epm	n/a	n/a	0.0	0.0	0.0	0.0	0.0	1.3	1.63
Slot-1	esrp	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.9	0.50
Slot-1	etmon	n/a	n/a	0.0	0.0	0.0	0.0	0.0	1.8	4.0
...										
...										
Slot-1	stp	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.9	0.67
Slot-1	telnetd	n/a	n/a	0.0	0.0	0.0	0.0	0.0	1.5	0.23
Slot-1	tftpd	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	0.19
Slot-1	thttpd	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.8	0.21
Slot-1	upm	n/a	n/a	0.0	0.0	0.0	0.0	0.0	1.8	0.43
Slot-1	vlan	n/a	n/a	0.0	0.0	0.0	0.0	0.1	4.3	4.28
Slot-1	vrrp	n/a	n/a	0.0	0.0	0.0	0.0	0.0	1.8	0.38
Slot-1	xmld	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	0.48
Slot-6	aaa	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	0.42
Slot-6	acl	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	0.40
Slot-6	bgp	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Slot-6	brm	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	0.18
Slot-6	cfgmgr	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	0.81
Slot-6	cli	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.9	7.17
Slot-6	devmgr	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	0.35
Slot-6	dirser	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Slot-6	dosprotect	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	0.6
Slot-6	eaps	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	0.60
Slot-6	edp	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	0.23
Slot-6	elrp	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.9	0.20
Slot-6	elsm	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	0.21
Slot-6	ems	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	0.44
Slot-6	epm	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	1.78
Slot-6	esrp	n/a	n/a	0.0	0.0	0.0	0.0	0.0	0.0	0.24
Slot-6	etmon	n/a	n/a	0.0	0.0	0.0	0.0	0.0	1.8	1.11
...										

History

This command was first available in an ExtremeXOS 11.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show database

```
show database database_name
```

Description

Shows information about all servers configured for Automation Edge remote VXLAN network identifier (VNI)-device databases.

Syntax Description

database	Shows database information.
<i>database_name</i>	Shows database information for the specific named database.

Default

N/A.

Usage Guidelines

N/A.

Example

The following example shows server information for all databases:

```
# show database
Retries : 3
-----
Database Name
  Server                               Port  Flags
-----
NewsroomNetwork1
  2001:0db8:85a3:0000:0000:8a2e:0370:7334  5456  C----
  CastleBlack                               34344  D----
  10.10.10.10                               12345  C----
NewsroomNetwork2
  3001:0db8:85a3:0000:0000:8a2e:0370:7334  5456  C----
  EastWatch                               34344  D----
  10.10.10.12                               12345  T----
-----
Flags : (C) Connected, (D) Disconnected, (T) Timed out.
```

The following example shows server information for the database "NewsroomNetwork1":

```
# show database NewsroomNetwork1
Database : NewsroomNetwork1                Retry : 3
  Server: 2001:0db8:85a3:0000:0000:8a2e:0370:7334 Port : 5456
  Status: DOWN
  Last connect : Tue Sep 10 20:22:48 UTC 2019
```



```
Last disconnect: Tue Sep 10 20:22:48 UTC 2019
# Disconnects   : 2
Server: EastWatch                               Port   : 34245
Status: UP
Last connect    : Tue Sep 10 20:22:48 UTC 2019
Last disconnect: Tue Sep 10 20:22:48 UTC 2019
# Disconnects   : 5
```

History

This command was first available in ExtremeXOS 31.1 as a demonstration feature.

This command is fully supported in ExtremeXOS 31.2.

Platform Availability

This command is available on the ExtremeSwitching 5520 platform.

show debug

```
show debug
```

Description

This command displays whether the writing of core dump files is enabled or disabled and whether the files are written to internal memory, a compact flash card, or a USB 2.0 storage device.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

By default, the switch writes core dump files to the internal memory card. To change the default configuration, use the `configure debug core-dumps [off | directory_path]` command.

Example

The following example shows that the switch is configured to send core dump files to the internal memory card:

```
Switch.2 # show debug
Debug Settings:
  Core dumps: Enabled (/usr/local/tmp)
```

The following example shows that the sending core dump files is disabled:

```
Switch.99 # show debug
Debug Settings:
  Core dumps: Disabled
```

The following example shows that the switch is configured to send core dump files to a compact flash card or USB 2.0 storage device:

```
Switch.76 # show debug
Debug Settings:
  Core dumps: Enabled (/usr/local/ext)
```

History

This command was first available in ExtremeXOS 11.1.

Support for the internal memory card was added in ExtremeXOS 11.2.

Support for USB 2.0 storage devices was added in ExtremeXOS 12.5.3.

The output was changed to show a directory path ExtremeXOS 30.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show dhcp-client state

```
show dhcp-client {ipv4 | ipv6} state {vlan}
```

Description

Displays the current *DHCP*/BOOTP client state for each vlan.

Syntax Description

This command has no arguments or variables.

Default

If the IPv4 or IPv6 keyword is not specified, IPv4 is the default.

Usage Guidelines

None.

Example

The following command displays the DHCP/BOOTP status for all VLANs:

```
show dhcp-client state
```

Depending on your configurations, output from this command is similar to the following:

```
# sh dhcp-client ipv6 state v3
Client VLAN           : v3
Client Type           : Stateful
Identity Association Type : IA-NA
                        IA-NA : Identity Association for Non-Temporary
Addresses
Identity Association Identifier : 96:7e:1f:5a
Client IP Address             : 2001::550b:ebf8:d0b8:9a8
Client Identifier Type        : Link-Layer Address Plus Time
Client Identifier             : 00:01:00:01:1a:85:0a:bd:00:04:96:7e:1a:5a:00:00
Server Identifier             : 00:01:00:01:13:e0:f4:93:00:11:43:c0:06:50
State                         : Bound

Identity Association Types:
  IA-NA - Identity Association for Non-Temporary Addresses,
  IA-TA- Identity Association for Temporary Addresses.
```

History

This command was first available in ExtremeXOS 10.1.

This command was modified in ExtremeXOS 15.6 to include the **ipv4** and **ipv6** keywords.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show dhcp-server

```
show dhcp-server {vlan vlan_name}
```

Description

Displays the *DHCP* server's configuration and address allocation on a specified *VLAN*.

Syntax Description

<i>vlan_name</i>	Specifies the VLAN of the DHCP server of interest.
------------------	--

Default

N/A.

Usage Guidelines

If no VLAN is specified, the configuration and address allocation for the servers on all the VLANs is displayed.

Example

The following command displays the configuration and address allocation for the DHCP server for the VLAN test:

```
show dhcp-server vlan test
```

The following is sample output from this command:

```
# show dhcp-server
VLAN "Default":
DHCP Address Range      : Not configured
Netlogin Lease Timer    : Not configured (Default = 10 seconds)
DHCP Lease Timer        : Not configured (Default = 7200 seconds)
DHCP Option Code  12    : hex "11:22:33:44:45"
DHCP Option Code  69    : ipaddress 10.0.0.1 10.0.0.2
Ports DHCP Enabled     : No ports enabled
```

History

This command was first available in ExtremeXOS 11.0.

The output is modified to show primary and secondary DNS servers in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show diagnostics

```
show diagnostics {slot [slot_number]}
```

Description

Displays the status of the last diagnostic test run on the switch.

Syntax Description

<i>slot</i>	Specifies which node to display diagnostic status information on.
<i>slot_number</i>	Specifies the slot number of the node on which you need to run the diagnostics. Values can be from 1 to 8.

Default

N/A.

Usage Guidelines

Use this command to display information from the last diagnostic test run on the switch.

To run diagnostics on a switch in a stack, you need to remove the node from the stacking mode, restart the node, and then run the diagnostics locally on the node. You can see the latest diagnostics results from the Master node of a stack after the node has rejoined the stack.

Output

The switch displays the following diagnostic information:

If you have run diagnostics, information includes the:

- Date the test was run—The month, date, and year.
- Last Test Version—The ExtremeXOS version associated with the results.
- Test data—If the diagnostic test failed, the switch displays the name of the failed test. The switch displays a message similar to the following: MAC memory test failed
- Summary—A brief summary of the overall diagnostic test. Options are:
 - Diagnostics Pass—The diagnostic test has passed.
 - Diagnostics Fail—One or more diagnostic tests have failed.

If you have never run diagnostics on the switch or stack ports, the switch displays a message similar to the following:

```
Result: FAIL
Test date run is invalid. Please run Diagnostics.
Error reading diagnostics information.
```

This message is normal and expected if you have never run diagnostics on the switch. After you run diagnostics, you should see information about the executed test.

Running Diagnostics

To run diagnostics, use the following command:

```
run diagnostics [extended | normal | stack-port]
```

Depending on the software version running on your switch or your switch model, additional or different diagnostic information might be displayed. For more information, see the command [run diagnostics](#).

Example

The following command displays the results of the switch diagnostics:

```
show diagnostics
```

The following is sample output from this command:

```
Last Test Date: May-04-2006
Last Test Version: 12.3.0.1
Summary: Diagnostics Pass
```

When the version is unknown, Last Test Version reads “Unknown.”

History

This command was available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show diffserv examination

```
show diffserv examination
```

Description

Displays the DiffServ-to-QoS profile mapping.

Syntax Description

N/A.

Default

N/A.

Usage Guidelines

Once you alter the default mappings, the “->” in the display (shown below) becomes “* >”.

Example

Because the SummitStack and ExtremeSwitching series switches have two default QoS profiles, you see different displays depending on the platform.

The following is sample output:

```
show diffserv examination
CodePoint->QOSProfile mapping:
00->QP1 01->QP1 02->QP1 03->QP1 04->QP1 05->QP1 06->QP1 07->QP1
08->QP1 09->QP1 10->QP1 11->QP1 12->QP1 13->QP1 14->QP1 15->QP1
16->QP1 17->QP1 18->QP1 19->QP1 20->QP1 21->QP1 22->QP1 23->QP1
24->QP1 25->QP1 26->QP1 27->QP1 28->QP1 29->QP1 30->QP1 31->QP1
32->QP1 33->QP1 34->QP1 35->QP1 36->QP1 37->QP1 38->QP1 39->QP1
40->QP1 41->QP1 42->QP1 43->QP1 44->QP1 45->QP1 46->QP1 47->QP1
48->QP1 49->QP1 50->QP1 51->QP1 52->QP1 53->QP1 54->QP1 55->QP1
56->QP8 57->QP8 58->QP8 59->QP8 60->QP8 61->QP8 62->QP8 63->QP8
```

History

This command was first available in ExtremeXOS 10.1.

The ports keyword was first available in ExtremeXOS 12.2.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show diffserv replacement

```
show diffserv replacement
```

Description

Displays the DiffServ replacement code-point values assigned to each [QoS](#) profile.

Syntax Description

N/A.

Default

N/A.

Usage Guidelines

Once you alter the default mappings, the “->” in the display (shown below) becomes “* >”.

Example

The following is sample output:

```
show diffserv replacement
QoSProfile->CodePoint mapping:
QP1->00
QP8->56
```

History

This command was first available in ExtremeXOS 10.1.

The ports keyword was first available in ExtremeXOS 12.2.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

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
```

Output from this command looks similar to the following:

```
Number of domain suffixes: 2
Domain Suffix 1:          njudah.local
Domain Suffix 2:          dbackman.com
Number of name servers: 2
Name Server 1:  172.17.1.104
Name Server 2:  172.17.1.123
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show dns cache analytics configuration

```
show dns cache analytics configuration {{vr} vr_name}
```


Description

Displays Domain Name System (DNS) analytics configuration.

Syntax Description

dns	Domain Network System.
cache	Specifies DNS cache.
analytics	Specifies showing DNS analytics information.
configuration	Specifies showing the DNS analytics configuration.
vr	Specifies the DNS analytics for a specific VR.
<i>vr_name</i>	Specifies the VR name. If not specified, the VR of the current command context is used.

Default

If not specified, the VR of the current command context is used.

Usage Guidelines

This command shows if DNS analytics is enabled and shows the configured maximum entries and timeout period.

Example

The following example shows the DNS analytics configuration for the current VR:

```
# show dns cache analytics configuration
DNS Cache Analytics : Enabled
Maximum Entries     : 10000
Timeout             : 1440
```

History

This command was first available in ExtremeXOS 30.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show dns cache analytics protected-client

```
show dns cache analytics protected-client [{vr} vr_name]
```

Description

Shows the Domain Name System (DNS) cache analytics protected client list for a virtual router (VR).

Syntax Description

dns	Domain Name System.
cache	Specifies showing DNS cache information.
analytics	Specifies showing DNS cache analytics information.
protected-client	Specifies showing the DNS cache analytics protected client list.
vr	Specifies the VR for the protected client list.
<i>vr_name</i>	Specifies the VR. If not specified, the VR of the current command context is used.

Default

If not specified, by default the VR of the current command context is used.

Usage Guidelines

None.

Example

The following example shows DNS analytics protected client list for VR "VR-Default":

```
# show dns cache analytics protected-client VR-Default
VR Name                               Client IP/Mask length
-----
VR-Default                           192.168.1.2/32
VR-Default                           192.168.3.0/32
```

History

This command was first available in ExtremeXOS 30.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show dns cache configuration

```
show dns cache configuration [{vlan} vlan_name | {vr} vr_name]
```

Description

Shows the Domain Name System (DNS) cache configuration (enabled or disabled) on a virtual router (VR) or VLAN.

Syntax Description

dns	Domain name system.
cache	Specifies showing the DNS cache information.
configuration	Specifies showing the DNS cache configuration (enabled or disabled).
vlan	Specifies showing the DNS cache configuration on a VLAN.
<i>vlan_name</i>	Specifies the VLAN name.
vr	Specifies showing the DNS configuration cache on a VR.
<i>vr_name</i>	Specifies the VR name. If not specified, the VR of the current command context is used.

Default

If no VR name is specified, the VR of the current command context is used.

Usage Guidelines

None.

Example

The following example shows the DNS cache configuration on VR "vrA":

```
# show dns cache configuration vrA
DNS Cache : Enabled on some VLANs of virtual router "vrA"
DNSSEC    : Enabled

VLAN                               DNS Cache Operational
-----
vlanA1                               Enabled    Up
```

History

This command was first available in ExtremeXOS 30.3.

Domain Name System Security Extension (DNSSEC) status was added in ExtremeXOS 31.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show dns cache analytics statistics

```
show dns cache analytics statistics {client client_ip domain
    domain_name } {detail} [{vr} vr_name}]
```

Description

Shows Domain Name System (DNS) cache analytics statistics for clients on a virtual router (VR).

Syntax Description

dns	Domain Name System.
cache	DNS cache.
analytics	Specifies showing information about the DNS cache analytics.
statistics	Specifies showing information about the DNS cache analytics statistics.
client	Specifies client statistics.
<i>client_ip</i>	Specifies which client by IP address.
domain	Specifies detailed domain.
<i>domain_name</i>	Specifies the domain.
detail	Displayed detailed statistics information.
vr	Specifies showing DNS cache analytics statistics on a VR.
<i>vr_name</i>	Specifies the VR name. If not specified, the VR of the current command context is used.

Default

If no VR name is specified, the VR of the current command context is used.

Usage Guidelines

This command displays DNS analytics statistics for a given VR. This command displays the total DNS queries received from clients, displaying IPv4 and IPv6 queries separately. For a specific client IP address, this command displays domain names queried by the client, number of DNS queries received for the domains, and the time stamp of the last received DNS query. Also, for a particular domain, it displays a detailed view of the specific domain with all the DNS query time stamps separately for IPv4 and IPv6.

Example

The following example shows DNS cache analytics statistics for all clients on the current VR:

```
# show dns cache analytics statistics
DNS Cache           : Enabled
DNS Cache Analytics : Enabled

Client IP                               # IPv4    # IPv6
                                   Queries    Queries
-----
10.127.2.251                           21928    2682
```

The following example shows DNS cache analytics statistics for the client at IP address "1.1.1.1" at domain "www.aaaaa.com" on the current VR:

```
# show dns cache analytics statistics client 1.1.1.1 domain www.aaaaa.com
DNS Cache           : Enabled
DNS Cache Analytics : Enabled

Domain Name          Query Type      Query Time
```

```

-----
www.aaaaa.com          IPv4 [A]          Thu Jan 15 06:56:07 1970
                        IPv4 [A]          Thu Jan 15 06:56:07 1970

Number of DNS IPv4 [A] Queries: 2
Number of DNS IPv6 [AAAA] Queries: 0

Flags: > - Indicates Domain Name is truncated past 44 characters.

```

The following example shows detailed DNS cache analytics statistics for the client at IP address "10.127.2.200" on the current VR:

```

# show dns cache analytics statistics client 10.127.2.200 detail
Domain Name           : geo.yahoo.com
# IPv4 Queries        : 1
# IPv6 Queries        : 0
Last IPv4 Query Time  : Fri Jan  4 06:48:56 2019
Last IPv6 Query Time  :

Domain Name           : acdn.adnxs.com
# IPv4 Queries        : 1
# IPv6 Queries        : 0
Last IPv4 Query Time  : Fri Jan  4 06:48:59 2019
Last IPv6 Query Time  :

Domain Name           : cdn.adnxs.com
# IPv4 Queries        : 1
# IPv6 Queries        : 0
Last IPv4 Query Time  : Fri Jan  4 06:48:59 2019
Last IPv6 Query Time  :

Domain Name           : sync.adaptv.advertising.com
# IPv4 Queries        : 1
# IPv6 Queries        : 0
Last IPv4 Query Time  : Fri Jan  4 06:49:02 2019
Last IPv6 Query Time  :

Domain Name           : ocsp.digicert.com
# IPv4 Queries        : 1
# IPv6 Queries        : 0
Last IPv4 Query Time  : Fri Jan  4 06:49:02 2019
Last IPv6 Query Time  :

Domain Name           : teredo.ipv6.microsoft.com
# IPv4 Queries        : 1
# IPv6 Queries        : 0
Last IPv4 Query Time  : Fri Jan  4 06:49:18 2019
Last IPv6 Query Time  :

```

The following example shows DNS cache analytics statistics for the domain "www.google.com" on the current VR:

```

# show dns cache analytics statistics domain www.google.com

IPv4      # IPv6      #
VR Name   Queries    Client IP
-----
VR-Default 1          192.168.1.1
Durga      1          192.168.2.1

```

```
2          0
```

History

This command was first available in ExtremeXOS 30.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show dns cache

```
show dns cache {current} {detail}
```

Description

Shows Domain Name System (DNS) cache entries.

Syntax Description

dns	Domain Name System.
cache	Specifies showing DNS cache entries.
current	Retrieves the current entries rather than most recent cache. Using this option might take some time to display the output.
detail	Displays output without truncation.

Default

N/A.

Usage Guidelines

To clear the DNS cache, use the command `clear dns cache`.

Example

The following example shows DNS cache entries (truncated):

```
# show dns cache
Domain Name                Address                Flag  Expiry Time
-----
oosp.starfieldtechaaaa.com>  oosp.godaddy.com.akadnssss.net>  CF    Fri Jan  4
07:32:45 2019
a652.dscb.akamai.net        182.156.239.25        4F    Fri Jan  4
06:35:47 2019
a652.dscb.akamai.net        182.156.239.27        4F    Fri Jan  4
06:35:47 2019
s.twitter.com               104.244.42.195        4F    Fri Jan  4
```

```

06:59:12 2019
s.twitter.com          104.244.42.67          4F    Fri Jan  4
06:59:12 2019
s.twitter.com          104.244.42.131        4F    Fri Jan  4
06:59:12 2019
s.twitter.com          104.244.42.3          4F    Fri Jan  4
06:59:12 2019
dmp.truoptik.com       104.16.91.60          4F    Fri Jan  4
06:37:35 2019
dmp.truoptik.com       104.16.92.60          4F    Fri Jan  4
06:37:35 2019

Flags: 4 - IPv4, 6 - IPv6, C - Canonical Name, F - Forwarded Query, R - Reverse
Query,
X - Non-Existent Domain Name, > - Indicates Domain Name or Address is truncated past 44
characters.
Total: 9

```

The following example shows DNS cache entries (not truncated):

```

# show dns cache detail
Domain Name : ojsp.starfieldtech.com
Address      : ojsp.godaddy.com.akadns.net
Flag        : CF
Expiry Time : Fri Jan  4 07:32:45 2019

Domain Name : a652.dscb.akamai.net
Address      : 182.156.239.25
Flag        : 4F
Expiry Time : Fri Jan  4 06:35:47 2019

Domain Name : a652.dscb.akamai.net
Address      : 182.156.239.27
Flag        : 4F
Expiry Time : Fri Jan  4 06:35:47 2019

Domain Name : s.twitter.com
Address      : 104.244.42.195
Flag        : 4F
Expiry Time : Fri Jan  4 06:59:12 2019

Domain Name : s.twitter.com
Address      : 104.244.42.67
Flag        : 4F
Expiry Time : Fri Jan  4 06:59:12 2019

Domain Name : s.twitter.com
Address      : 104.244.42.131
Flag        : 4F
Expiry Time : Fri Jan  4 06:59:12 2019

Flags: 4 - IPv4, 6 - IPv6, C - Canonical Name, F - Forwarded Query, R - Reverse Query,
      X - Non-Existent Domain Name
Total: 9

```

History

This command was first available in ExtremeXOS 30.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show dns cache name-server

```
show dns cache name-server
```

Description

Displays the Domain Name System (DNS) name server.

Syntax Description

dns	Domain Name System.
cache	Specifies viewing DNS name servers.
name-server	Specifies viewing DNS name servers.

Default

N/A.

Usage Guidelines

None.

Example

The following example shows DNS name servers:

```
# show dns cache name-server
Name Server 1: 1.1.1.2          VR: VR-Default
```

History

This command was first available in ExtremeXOS 30.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show dos-protect

```
show dos-protect {detail}
```

Description

Displays DoS protection configuration and state.

Syntax Description

detail	Specifies to display statistics in addition to configuration and state.
---------------	---

Default

N/A.

Usage Guidelines

Use this command to display the DoS protection settings. Using the detail option will also display the following cumulative statistics:

- trusted.
- notify.
- alerts.

Example

The following command displays the DoS protection settings for the switch:

```
show dos-protect
```

The following is sample output from this command:

```
dos-protect is disabled
dos-protect settings:
interval:          1  (measurement interval secs)
acl expire time:   5  (secs)
trusted ports:
no trusted ports configured
type L3-Protect:
notify threshold:   3500  (level to log a message)
alert threshold:    4000  (level to generate an ACL)
```

The following command displays detailed DoS protection settings for the switch:

```
show dos-protect detail
```

The following is sample output from this command:

```
dos-protect is enabled
dos-protect settings:
interval:          1  (measurement interval secs)
acl expire time:   5  (secs)
trusted ports:
1:2
type L3-Protect:
notify threshold:   3500  (level to log a message)
alert threshold:    4000  (level to generate an ACL)
dos-protect statistics:
trusted:    1301
notify:     0
alerts:     0
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show dot1p

show dot1p

Description

Displays the 802.1p-to-QoS profile mappings.

Syntax Description

N/A.

Default

N/A.

Usage Guidelines

N/A.

Example

The SummitStack and ExtremeSwitching series switches have two default [QoS](#) profiles.

Following is sample output from the show dot1p command:

```
show dot1p
802.1p Priority Value      QOS Profile      Ingress Meter
      0                QP1              ingmeter0
      1                QP1              none
      2                QP1              ingmeter0
      3                QP1              ingmeter4
      4                QP1              ingmeter4
      5                QP1              ingmeter15
      6                QP1              none
      7                QP8              ingmeter8
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show dwdm channel-map

```
show dwdm channel-map { channel_first { - channel_last } } {port
    port_num}
```

Description

Displays the channel scheme adopted for mapping the DWDM wavelengths.

Syntax Description

<i>channel_first</i>	Specifies the starting channel number.
<i>channel_last</i>	Specifies the ending channel number.
<i>port_num</i>	Specifies the port for which the status is to be displayed.

Default

N/A.

Usage Guidelines

Use this command to display the wavelength and the supportability of the channel by the optical module in the port.

Example

The following command displays information for channels 50 through 60 on port 3:1:

```
show dwdm channel-map 50 - 60 port 3:1
```

The following is sample output for this command:

```
=====
Channel #      Wavelength (nm)      Port 3:1
=====
50             1537.40             Supported
51             1536.61             Supported
52             1535.82             Supported
...            .....
58             1531.12             Supported
59             1530.33             Supported
60             1529.55             Supported
```

History

This command was first available in ExtremeXOS 12.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show eaps

```
show eaps {eapsDomain} {detail}
```

Description

Displays *EAPS* status information.

Syntax Description

<i>eapsDomain</i>	Specifies the name of an EAPS domain.
detail	Specifies all available detail for each domain.

Default

N/A.

Usage Guidelines

If you enter the show eaps command without a keyword, the command displays less than with the detail keyword.

Use the optional eapsDomain parameter to display status information for a specific EAPS domain.

Some state values are different on a transit node than on a master node.

When you enter the show eaps command without a domain name, the switch displays the following fields:

EAPS Enabled:	Current state of EAPS on this switch: Yes—EAPS is enabled on the switch.No—EAPS is not enabled.
EAPS Fast Convergence:	Displays only when Fast Convergence is on.
EAPS Display Config Warnings:	Displays the setting for loop protection messages: On—Loop protection messages are displayed (this is the default behavior).Off—Loop protection messages are not displayed.
EAPS Multicast Add Ring Ports:	Displays the configuration of the multicast add-ring-ports feature as configured with the <code>configure eaps multicast add-ring-ports</code> command.
EAPS Multicast Send <i>IGMP</i> Query:	Displays the configuration of the multicast send-igmp-query feature as configured with the <code>configure eaps multicast send-igmp-query</code> command.
EAPS Multicast Temporary Flooding:	Displays the configuration of the multicast temporary-flooding feature as configured with the <code>configure eaps multicast temporary-flooding</code> command.

EAPS Multicast Temporary Flooding Duration:	Displays the duration configuration for the multicast temporary-flooding feature as configured with the <code>configure eaps multicast temporary-flooding duration</code> command.
Number of EAPS instances:	Number of EAPS domains created. The maximum number of EAPS domains per switch is 128.
Domain:	Entries in this column identify the name of an EAPS domain.
State:	<p>On a transit node, the command displays one of the following states:</p> <p>Idle—The EAPS domain has been enabled, but the configuration is not complete.</p> <p>Links-Up—This EAPS domain is running, and both its ports are up and in the FORWARDING state.</p> <p>Links-Down—This EAPS domain is running, but one or both of its ports are down.</p> <p>Preforwarding—This EAPS domain is running, and both of its ports are up, but one of them is in a temporary BLOCKED state.</p> <p>On a master node, the command displays one of the following states:</p> <p>Idle—The EAPS domain has been enabled, but the configuration is not complete.</p> <p>Init—The EAPS domain has started but has not yet determined the status of the ring. The secondary port is in a BLOCKED state.</p> <p>Complete—The ring is in the COMPLETE state for this EAPS domain.</p> <p>Failed—There is a break in the ring for this EAPS domain.</p> <p>Pre-Init—The EAPS domain has started operation for Init state and has sent a request to lower hardware layers to block the secondary port. It is in transient state waiting for acknowledgement from hardware layer indicating the operation is completed.</p> <p>Pre-Complete—The EAPS domain has started operation for Complete state and has sent a request to lower hardware layers to block the secondary port. It is in transient state waiting for acknowledgement from the hardware layer indicating the operation is completed.</p> <p>[Failtimer Expired]—When the failtimer expires and its action is set to send-alert, this flag is set. This flag indicates there is a misconfiguration or hardware problem in the EAPS ring. The EAPS master node continues to remain in COMPLETE or INIT state with its secondary port blocking.</p>
Mo:	The configured EAPS mode for this switch: transit (T) or master (M).
Primary/Secondary port:	The port numbers assigned as the EAPS primary and secondary ports. On the master node, the port distinction indicates which port is blocked to avoid a loop.
Prio	The EAPS domain priority, which is H for high priority or N for normal priority.

When you enter the `show eaps` command with a domain name or the detail keyword, the switch displays the following fields:

Name:	Identifies the EAPS domain displayed.
Priority	The EAPS domain priority, which is either High or Normal.

State:	<p>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.</p> <p>On a master node, the command displays one of the following states: Idle—The EAPS domain has been enabled, but the configuration is not complete. Init—The EAPS domain has started but has not yet determined the status of the ring. The secondary port is in a BLOCKED state. Complete—The ring is in the COMPLETE state for this EAPS domain. Failed—There is a break in the ring for this EAPS domain. Pre-Init—The EAPS domain has started operation for Init state and has sent a request to lower hardware layers to block the secondary port. It is in transient state waiting for acknowledgement from hardware layer indicating the operation is completed. Pre-Complete—The EAPS domain has started operation for Complete state and has sent a request to lower hardware layers to block the secondary port. It is in transient state waiting for acknowledgement from the hardware layer indicating the operation is completed. [Failtimer Expired]—When the failtimer expires and it's action is set to send-alert, this flag is set. This flag indicates there is a misconfiguration or hardware problem in the EAPS ring. The EAPS master node continues to remain in COMPLETE or INIT state with it's secondary port blocking.</p>
[Running: ...]	Yes—This EAPS domain is running. No—This EAPS domain is not running.
Enabled:	Indicates whether EAPS is enabled on this domain. Y—EAPS is enabled on this domain. N—EAPS is not enabled.
Mode:	The configured EAPS mode for this switch: transit (T) or master (M).
Primary/Secondary port:	The port numbers assigned as the EAPS primary and secondary ports. On the master node, the port distinction indicates which port is blocked to avoid a loop.
Port status:	Unknown—This EAPS domain is not running, so the port status has not yet been determined. Up—The port is up and is forwarding data. Down—The port is down. Blocked—The port is up, but data is blocked from being forwarded.
Tagstatus:	Tagged status of the control <u>VLAN</u> : Tagged—The control VLAN has this port assigned to it, and the port is tagged in the VLAN. Untagged—The control VLAN has this port assigned to it, but the port is untagged in the control VLAN. Undetermined—Either a VLAN has not been added as the control VLAN to this EAPS domain or this port has not been added to the control VLAN.
Hello timer interval:	The configured value of the timer in seconds and milliseconds, specifying the time that the master node waits between transmissions of health check packets.
Fail timer interval:	The configured value of the timer in seconds, specifying the time that the master node waits before the failtimer expires.

Failtimer expiry action:	Displays the action taken when the failtimer expires: Send-alert—Sends a critical message to the syslog when the failtimer expires. Open-secondary-port—Opens the secondary port when the failtimer expires. Displays only for master nodes.
Preforwarding Timer interval:	The configured value of the timer. This value is set internally by the EAPS software. The set value is 15 seconds. Note: If two links in an EAPS domain go down at the same time and one link comes back up, it takes 15 seconds for the reconnected link to start receiving traffic again. Displays only for transit nodes.
Last valid EAPS update:	Indicates the last time a hello packet was received.
EAPS Domain Controller Vlan:	Lists the assigned name and ID of the control VLAN.
EAPS Domain Protected Vlan(s):	Lists the assigned names and VLAN IDs of all the protected VLANs configured on this EAPS domain.
Number of Protected Vlans:	The count of protected VLANs configured on this EAPS domain.

Example

The following command displays information for all EAPS domains:

```
Switch.5 # show eaps
EAPS Enabled: Yes
EAPS Fast-Convergence: Off
EAPS Display Config Warnings: On
EAPS Multicast Add Ring Ports: Off
EAPS Multicast Send IGMP Query: On
EAPS Multicast Temporary Flooding: Off
EAPS Multicast Temporary Flooding Duration: 15 sec
Number of EAPS instances: 2
# EAPS domain configuration :
-----
Domain          State          Mo  En  Pri  Sec  Control-Vlan VID  Count Prio
-----
d1              Idle          T   N   1    2   cv1              (101 ) 0    H
d2              Links-Up      T   Y  3:8  3:16 c2              (1001) 100  H
-----
```

The following command displays information for EAPS domain d1:

```
Switch.7 # show eaps d1
Name: d1          Priority: High
State: Idle       Running: No
Enabled: No       Mode: Transit
Primary port: 1    Port status: Unknown  Tag status: Undetermined
Secondary port: 2  Port status: Unknown  Tag status: Undetermined
Hello timer interval: 1 sec 0 millsec
Fail timer interval: 3 sec 0 millsec
Fail Timer expiry action: Send alert
Last valid EAPS update: From Master Id 00:01:30:f9:9c:b0, at Wed Jun 9 09:09:35 2004
EAPS Domain has following Controller Vlan:
Vlan Name      VID
c1             1000
```

³ These fields apply only to transit nodes; they are not displayed for a master node.

```

EAPS Domain has following Protected Vlan(s):
Vlan Name          VID
p_1                 1
p_2                 2
p_3                 3
p_4                 4
p_5                 5
p_6                 6
p_7                 7
p_8                 8
p_9                 9
p_10                10
p_11                11
p_12                12
p_13                13
p_14                14
p_15                15
p_16                16
p_17                17
p_18                18
p_19                19
p_20                20
p_21                21
p_22                22
p_23                23
p_24                24
p_25                25
p_26                26
p_27                27
p_28                28
p_29                29
p_30                30

```

The following command displays information on EAPS domain domain12, which is configured to send hello packets on the secondary port:

```

Switch.9 # show eaps "domain12"
Name: domain12      Priority: High
State: Complete     Running: Yes
Enabled: Yes        Mode: Master
Primary port: 17     Port status: Up Tag status: Tagged
Secondary port: 27   Port status: Blocked Tag status: Tagged
Hello Egress Port: Secondary
Hello timer interval: 0 sec 100 millisec
Fail timer interval: 0 sec 300 millisec
Fail Timer expiry action: Send alert
Last update: From Master Id 00:04:96:34:e3:43, at Tue May 11 15:39:29 2010
EAPS Domain has following Controller Vlan:
Vlan Name          VID
vlancl2            1002
EAPS Domain has following Protected Vlan(s):
Vlan Name          VID
pvlan11            204
pvlan12            205
pvlan13            206
Number of Protected Vlans: 3

```



Note

You might see a slightly different display, depending on whether you display the master node or the transit node.

The display from the `show eaps detail` command shows all the information shown in the `show eaps eapsDomain` command, but displays information for all configured EAPS domains.

For the CFM support in EAPS, the existing `show eaps` output places a “!” next to a CFM monitored ring port if the CFM indicates the MEP group for that port is down.

```
# sh eaps
EAPS Enabled: Yes
EAPS Fast-Convergence: Off
EAPS Display Config Warnings: Off
EAPS Multicast Add Ring Ports: Off
EAPS Multicast Send IGMP Query: On
EAPS Multicast Temporary Flooding: Off
EAPS Multicast Temporary Flooding Duration: 15 sec
Number of EAPS instances: 1
# EAPS domain configuration :
-----
Domain          State          Mo  En  Pri   Sec   Control-Vlan VID   Count Prio
-----
d2              Failed          M   Y   !41    31    v2              (101 ) 1    N
-----
Flags : (!) CFM Down
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show eaps cfm groups

```
show eaps cfm groups
```

Description

Displays summary *EAPS* CFM groups information.

Syntax Description

There are no keywords or variables for this command.

Default

N/A.

Usage Guidelines

The following command displays EAPS CFM group information:

```
# sh eaps cfm groups
-----
MEP Group Name          Status Port   MEP ID
-----
eapsCfmGrp1             Up    41        11
eapsCfmGrp2             Up    31        12
```

History

This command was first available in ExtremeXOS 15.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show eaps counters shared-port

```
show eaps counters shared-port [global | port {segment-port segport
    {eapsDomain}}]
```

Description

Displays summary EAPS shared port counter information.

Syntax Description

global	Displays general counter information for all configured EAPS shared port instances. The output displayed is calculated for all configured EAPS shared ports; not just one specific shared port instance.
<i>port</i>	Identifies the port number of the specified common link port.
<i>segport</i>	Identifies the segment port. The segment port is the other ring port of an EAPS domain that is not the shared-port.
<i>eapsDomain</i>	Specifies the name of the EAPS domain. If no EAPS domain is specified, all counters for all EAPS domains on the specified segment port are displayed.

Default

N/A.

Usage Guidelines

If the switch is configured for EAPS shared ports, use this command to display an array of counters associated with the EAPS shared port functionality.

If you specify the global keyword, the switch displays general counter information for all configured EAPS shared port instances. The output displayed is calculated for all configured EAPS shared ports; not just one specific shared port instance.

If you specify a particular EAPS shared port, the switch displays counter information related to only that shared port.

If you specify a particular EAPS segment port, the switch displays counter information related to only that segment port for the specified EAPS domain.

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.

Clearing the Counters

The counters continue to increment until you clear the information. By clearing the counters, you can see fresh statistics for the time period you are monitoring. To clear, reset the EAPS counters, including the shared port counters, use one of the following commands:

- `clear counters`
- `clear eaps counters`

Understanding the Output

The following table describes the significant fields and values in the display output of the `show eaps counters shared-port` global command:

Field	Description
Rx-Invalid-Instance	Displays the number of dropped EAPS shared-port PDUs because there is not a valid EAPS shared port instance for the incoming port.
Rx-Unknown	Displays the number of unknown EAPS PDUs dropped by the shared port instances.
Fw-Invalid-Instance	Displays the number of EAPS shared-port PDUs that could not be forwarded in slow path because the shared port instances could not find a valid EAPS shared port instance for the outgoing port.

The following table describes the significant fields and values in the display output of the `show eaps counters shared-port portsegment-port segport eapsDomain` command:

Field	Description
Rx-Seg-Health	Indicates the shared port instance received EAPS shared ports Segment-Health-Check PDUs.
Rx-Path-Detect	Indicates the shared port instance received EAPS shared ports Path-Detect PDUs.

Field	Description
Rx-Flush-Notify	Indicates the shared port instance received EAPS shared ports Flush-Notify PDUs and flushed the <i>FDB</i> . If this PDU reaches a port of the shared ports pair that initiated the PDU, the shared port instance might terminate the PDU. Otherwise, the shared port instance forwards the PDU.
Rx-Unknown	Displays the number of unknown EAPS PDUs dropped by the shared port instance.
Rx-Seg-Health-Dropped	Displays the number of EAPS shared ports Segment-Health-Check PDUs dropped by the shared port instance. This counter increments if the Segment-Health-Check PDU returns to the sending switch. If that occurs, the switch drops the Segment-Health-Check PDU.
Rx-Path-Detect-Dropped	Displays the number of EAPS shared ports Path-Detect PDUs dropped by the shared port instance. This counter increments in the following situations: If the packet's Fwd-id matches the EAPS shared port's Link-Id, the port is not in the blocking state, and the incoming port is a segment port. If the packet's Link-Id matches the EAPS shared port's Link-Id, the port is not in the blocking state, and the incoming port is a segment port.
Rx-Flush-Notify-Dropped	Displays the number of EAPS shared ports Flush-Notify-Dropped PDUs dropped by the shared port instance. This counter increments in the following situations: If the Flush-Notify-Dropped PDU returns to the sending switch. If the packet's Fwd-Id matches the EAPS shared port's Link-Id and the port is not in the blocking state.
Rx-Dropped-Invalid-Port	Displays the number of EAPS shared ports PDUs dropped by the shared port instance because it does not exist.
Tx-Seg-Health	Indicates the shared port instance sent EAPS shared ports Segment-Health-Check PDUs.
Tx-Path-Detect	Indicates the shared port instance sent EAPS shared ports Path-Detect PDUs. NOTE: This counter appears under Common Link Port Stats and should always be 0.
Tx-Flush-Notify	Indicates the shared port instance sent EAPS shared ports Flush-Notify PDUs to flush the FDB. NOTE: This counter appears under Common Link Port Stats and should always be 0.
Tx-Flush-Fdb	Indicates the shared port instance sent EAPS Flush-Fdb PDUs because the FDB needs to be flushed. NOTE: This counter appears under Common Link Port Stats and should always be 0.
Tx-Unknown	Indicates the number of unknown EAPS PDUs sent by the shared port instance. NOTE: Unknown EAPS PDUs can be a new type of PDU that the switch does not track in the sending routine.
Tx-Transmit-Err	Indicates the number of EAPS PDUs the shared port instance was unable to send because of an error.
Fw-Seg-Health	Indicates the number of EAPS shared ports Segment-Health-Check PDUs received by the shared port instance and forwarded in slow path.

Field	Description
Fw-Path-Detect	Indicates the number of EAPS shared ports Path-Detect PDUs received by the shared port instance and forwarded in slow path.
Fw-Flush-Notify	Indicates the number of EAPS Flush-Notify PDUs received by the shared port instance and forwarded in slow path to flush the FDB.
Fw-Flush-Fdb	Indicates the number of EAPS Flush-Fdb PDUs received by the shared port instance and forwarded in slow path.
Fw-Unknown	Indicates the number of unknown EAPS PDUs forwarded in slow path. NOTE: Unknown EAPS PDUs can be a new type of PDU that the switch does not track in the forwarding routine.
Fw-Transmit-Err	Indicates the number of EAPS PDUs the shared port instance was unable to forward in slow path because of an error.

Example

The following command displays global, high-level counter information for EAPS shared port:

```
show eaps counters shared-port global
```

The following is sample output from this command:

```
Global counters for EAPS Shared-Ports:
Rx Dropped
Rx-Invalid-Instance : 0
Rx-Unknown          : 0
Fw Dropped
Fw-Invalid-Instance : 0
```

The following example assumes that port 17 is configured as an EAPS shared port. The following command displays counter information the specified EAPS shared port:

```
show eaps counters shared-port 17
```

The following is sample output from this command:

```
Counters for EAPS Shared-Port 17:
Common Link Port Stats
Rx Stats
Rx-Seg-Health          : 0
Rx-Path-Detect         : 0
Rx-Flush-Notify        : 0
Rx Dropped
Rx-Seg-Health-Dropped  : 0
Rx-Path-Detect-Dropped : 0
Rx-Flush-Notify-Dropped : 0
Rx-Dropped-Invalid-Port : 0
Tx Stats
Tx-Seg-Health          : 0
Tx-Path-Detect         : 0
Tx-Flush-Notify        : 0
Tx-Flush-Fdb           : 0
Tx Dropped
Tx-Unknown             : 0
```

```

Tx-Transmit-Err      : 0
Fw Stats
Fw-Seg-Health        : 0
Fw-Path-Detect       : 0
Fw-Flush-Notify      : 0
Fw Dropped
Fw-Unknown           : 0
Fw-Transmit-Err      : 0

```

The following example assumes that port 1:2 is configured as an EAPS shared port and port 1:1 is a segment port. The following command displays counter information the specified EAPS shared port, segment port, and EAPS domain:

```
show eaps counters shared-port 1:2 segment-port 1:1 eaps1
```

The following is sample output from this command:

```

Counters for EAPS Shared-Port 1:2, Segment Port: 1:1, EAPS Domain: eaps1
Rx Stats
Rx-Seg-Health          : 0
Rx-Path-Detect         : 0
Rx-Flush-Notify        : 0
Rx-Seg-Health-Dropped  : 0
Rx-Path-Detect-Dropped : 0
Rx-Flush-Notify-Dropped : 0
Rx-Dropped-Invalid-Port : 0
Tx Stats
Tx-Seg-Health          : 2275
Tx-Path-Detect         : 0
Tx-Flush-Notify        : 0
Tx-Flush-Fdb           : 0
Tx-Transmit-Err        : 0
Tx-Unknown             : 0
Fw Stats
Fw-Seg-Health          : 0
Fw-Path-Detect         : 0
Fw-Flush-Notify        : 0
Fw-Transmit-Err        : 0
Fw-Unknown             : 0

```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on all platforms with the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and what licenses are appropriate for this feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

show eaps counters

```
show eaps counters [eapsDomain | global]
```

Description

Displays summary *EAPS* counter information.

Syntax Description

<i>eapsDomain</i>	Specifies the name of an EAPS domain. The switch displays counter information for only that domain.
global	Displays EAPS counter information when the events counted are not applicable to any specific EAPS domain.

Default

N/A.

Usage Guidelines

If you specify the name of an EAPS domain, the switch displays counter information related to only that domain. If you specify the global keyword, the switch displays EAPS counter information when the events counted are not applicable to any specific EAPS domain. The output displayed is for all configured EAPS domains, not just one specific EAPS domain.

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.

Clearing the Counters

The counters continue to increment until you clear the information. By clearing the counters, you can see fresh statistics for the time period you are monitoring. To clear, reset the EAPS counters, use one of the following commands:

- `clear counters`
- `clear eaps counters`

Understanding the Output

The following table describes the significant fields and values in the display output of the `show eaps counters eapsDomain` command:

Field	Description
Rx-Health	Indicates the EAPS domain received EAPS Health PDUs.
Rx-RingUp-FlushFdb	Indicates the EAPS ring is up, and the EAPS domain received EAPS RingUp-FlushFdb PDUs to flush the <i>FDB</i> .
Rx-RingDown-FlushFdb	Indicates the EAPS ring is down, and the EAPS domain received EAPS RingDown-FlushFdb PDUs to flush the FDB.
Rx-Link-Down	Indicates the EAPS domain received EAPS Link-Down PDUs and took down the link.

Field	Description
Rx-Flush-Fdb	Indicates the EAPS domain received EAPS Flush-Fdb PDUs and flushed the FDB.
Rx-Suspend-Prefwd-Timer	Indicates the EAPS domain received EAPS Suspend-Preforward-Timer PDUs.
Rx-Query-Link-Status	Indicates the EAPS domain received EAPS Query-Link-Status PDUs.
Rx-Link-Up	Indicates the EAPS domain received EAPS Link-Up PDUs and brought the link back up.
Rx-Unknown	Indicates the EAPS domain dropped unknown EAPS PDUs.
Rx-Another-Master	Indicates the EAPS domain dropped EAPS PDUs because there is another Master switch in the same EAPS domain.
Rx-Unconfigured-Port	Indicates the EAPS domain dropped EAPS PDUs because the ingress port is not configured to be a ring port for the EAPS domain and the corresponding control <u>VLAN</u> .
Rx-Health-Pdu-Pri-Port	Indicates the EAPS domain dropped EAPS Health PDUs because the primary port received them instead of the secondary port. NOTE: The secondary port of the Master switch must receive EAPS Health PDUs, not the primary port.
Tx-Health	Indicates the EAPS domain sent EAPS Health PDUs.
Tx-RingUp-FlushFdb	Indicates the EAPS ring is up, and the EAPS domain sent EAPS RingUp-FlushFdb PDUs to flush the FDB.
Tx-RingDown-FlushFdb	Indicates the EAPS ring is down, and the EAPS domain sent EAPS RingDown-FlushFdb PDUs to flush the FDB.
Tx-Link-Down	Indicates the EAPS domain sent EAPS Link-Down PDUs because the link went down.
Tx-Flush-Fdb	Indicates the EAPS domain sent EAPS Flush-Fdb PDUs because the FDB needs to be flushed.
Tx-Suspend-Prefwd-Timer	Indicates the EAPS domain sent EAPS Suspend-Preforward-Timer PDUs.
Tx-Query-Link-Status	Indicates the EAPS domain sent EAPS Query-Link-Status PDUs.
Tx-Link-Up	Indicates the EAPS domain sent EAPS Link-Up PDUs and the link is up.
Tx-Unknown	Indicates the number of unknown EAPS PDUs sent by the EAPS domain. NOTE: Unknown EAPS PDUs can be a new type of PDU that the switch does not track in the sending routine.
Tx-Transmit-Err	Indicates the number of EAPS PDUs the EAPS domain was unable to send because of an error.
Fw-Link-Down	Indicates the number of EAPS Link-Down PDUs received by the EAPS domain and forwarded in slow path.
Fw-Flush-Fdb	Indicates the number of EAPS Flush-Fdb PDUs received by the EAPS domain and forwarded in slow path.
FW-Query-Link-Status	Indicates the number of EAPS Query-Link-Status PDUs received by the EAPS domain and forwarded in slow path.

Field	Description
Fw-Unknown	Indicates the number of unknown EAPS PDUs forwarded in slow path. NOTE: Unknown EAPS PDUs can be a new type of PDU that the switch does not track in the forwarding routine.
Fw-Transmit-Er	Indicates the number of EAPS PDUs the EAPS domain was unable to forward in slow path because of an error.

**Note**

Rx and Fw counters—If a PDU is received, processed, and consumed, only the Rx counter increments. If a PDU is forwarded in slow path, both the Rx counter and Fw counter increment.

The following table describes the significant fields and values in the display output of the `show eaps counters` global command:

Field	Description
Rx-Failed	Indicates an error occurred when receiving packets from the Layer 2 forwarding engine.
Rx-Invalid-Vlan-Intf	Indicates that the VLAN interface for the incoming VLAN cannot be found.
Rx-Undersize-Pkt	Indicates the length of the packet is less than the length of the header.
Rx-Invalid-8021Q-Tag	Indicates the VlanTypeLength field in the Ethernet header does not match the default Ethernet value for the 802.1Q tag.
Rx-Invalid-SNAP-Type	Indicates an invalid Subnetwork Access Protocol (SNAP) value in the Ethernet header.
Rx-Invalid-OUI	Indicates the Organizational Unique Identifier (OUI) value in the Ethernet header does not match 00:E0:2B.
Rx-EEP-Unsupported-Version	Indicates an unsupported Extreme Encapsulation Protocol (EEP) version. The EEP version should be 1.
Rx-EEP-Invalid-Length	Indicates the length of the EEP header is greater than the length of the packet.
Rx-EEP-Checksum-Invalid	Indicates the EEP checksum is invalid.
Rx-Domain-Invalid	Indicates the control VLAN's incoming PDU is not associated with an EAPS domain.
Rx-Lif-Invalid	Indicates that EAPS is unable to determine the logical interface (LIF) for the ingress port.
Rx-Lif-Down	Indicates the LIF for the ingress port is in the Down state.
Tx-Failed	Indicates an error occurred when sending packets to the Layer 2 forwarding engine.

Example

The following command displays the counters for a specific EAPS domain named eaps1:

```
show eaps counters eaps1
```

The following is sample output from this command:

```
Counters for EAPS domain: eaps1
Rx Stats
Rx-Health                : 0
Rx-Ringup-Flushfdb       : 0
Rx-Ringdown-Flushfdb     : 0
Rx-Link-Down             : 0
Rx-Flush-Fdb             : 0
Rx-Suspend-Prefwd-Timer  : 0
Rx-Query-Link-Status     : 0
Rx-Link-Up               : 0
Rx Dropped
Rx-Unknown               : 0
Rx-Another-Master        : 0
Rx-Unconfigured-Port     : 0
Rx-Health-Pdu-Pri-Port   : 0
Tx Stats
Tx-Health                : 5011
Tx-Ringup-Flushfdb       : 0
Tx-Ringdown-Flushfdb     : 0
Tx-Link-Down             : 0
Tx-Flush-Fdb             : 0
Tx-Suspend-Prefwd-Timer  : 0
Tx-Query-Link-Status     : 3342
Tx-Link-Up               : 0
Tx Dropped
Tx-Unknown               : 0
Tx-Transmit-Err          : 0
Fw Stats
Fw-Link-Down             : 0
Fw-Flush-Fdb             : 0
Fw-Query-Link-Status     : 0
Fw Dropped
Fw-Unknown               : 0
Fw-Transmit-Err          : 0
```

The following command displays the global EAPS counters:

```
show eaps counters global
```

The following is sample output from this command:

```
Global counters for EAPS:
Rx-Failed : 0
Rx-Invalid-Vlan-Intf : 0
Rx-Undersize-Pkt : 0
Rx-Invalid-SNAP-Type : 0
Rx-Invalid-OUI : 0
Rx-EEP-Unsupported-Version : 0
Rx-EEP-Invalid-Length : 0
Rx-EEP-Checksum-Invalid : 0
Rx-Domain-Invalid : 0
Rx-Failed : 0
```

```
Rx-Lif-Invalid : 0
Rx-Lif-Down   : 0
Tx-Failed     : 0
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show eaps shared-port

```
show eaps shared-port {port} {detail}
```

Description

Displays shared-port information for one or more [EAPS](#) domains.

Syntax Description

<i>port</i>	Specifies a shared-port.
detail	Specifies to display the status of all segments and VLANs.

Default

N/A.

Usage Guidelines

If you enter the `show eaps shared-port` command without the `detail` keyword, the command displays a summary of status information for all configured EAPS shared ports.

If you specify an EAPS shared-port, the command displays information about that specific port and the related segment ports. The segment ports are sorted in ascending order based on their port number. You can use this order and your knowledge of the EAPS topology to determine which segment port becomes the active-open port if the common link fails. For more information, see “Common Link Fault Detection and Response” in the [Switch Engine 32.3 User Guide](#).

You can use the `detail` keyword to display more detailed status information about the segments and VLANs associated with each shared port.

The following table describes the significant fields and values in the display output of the `show eaps shared-port {port {detail}}` commands:

Field	Description
Shared Port	Displays the port number of the shared port.
Mode	Indicates whether the switch on either end of the common link is a controller or partner. The mode is configured by the user.
Link ID	The link ID is the unique common link identifier configured by the user.
Up	Displays one of the following: Yes—Indicates that the link ID and the mode are configured. No—Indicates that the link ID or the mode is not configured.
State	Displays one of the following states: Idle—Shared-port instance is not running. Ready—The EAPS shared-port instance is running, the neighbor can be reached, and the common link is up. Blocking—The EAPS shared-port instance is running, the neighbor cannot be reached, or the common link is down. Preforwarding—The EAPS shared-port instance is in a blocking state, and the common link came up. To prevent a superloop, a temporary blocking state is created before going into Ready state.
Domain Count	Indicates the number of EAPS domains sharing the common link.
VLAN Count	Indicates the total number of VLANs that are protected under the EAPS domains sharing this common link.
Nbr	Yes—Indicates that the EAPS instance on the other end of the common link is configured with matching link ID and opposite modes. For example, if one end of the common link is configured as a controller, the other end must be configured as a partner. Err—Indicates that the EAPS instance on the other end of the common link is configured with a matching link ID, but the modes are configured the same. For example, both modes are configured as controller, or both modes are configured as partner. No—The neighbor on the other end of the common link cannot be reached. Indicates one or more of the following: <ul style="list-style-type: none"> The switch on the other end of the common link is not running. The shared port has not been created. The link IDs on each side of the common link do not match. The common link, and any other segment, between the controller and partner are not fully connected.
RB ID	The ID of the root blocker. If the value is none, there are not two or more common-link failures.
RB State	None—This EAPS shared-port is not the root blocker. Active—This EAPS shared-port is the root blocker and is currently active. Inactive—This EAPS shared-port is the root blocker but is currently inactive.
Active Open (available with the detail keyword)	None—Indicates that there is no Active-Open port on the VLAN. Port #—Indicates the port that is Active-Open and is in a forwarding state.

Field	Description
Segment Timer expiry action	<p>Segment down—Specifies that if the controller or partner switch detects a down segment, that segment stays down and a query is not sent through the ring. The switch marks the segment status as Down.</p> <p>Send alert—Specifies that if the controller or partner switch detects a down segment, that switch keeps the segment up and sends a warning message to the log (default). The switch sends a trap alert and sets the failed flag [F].</p>
Segment Port (available with the detail keyword or by specifying a shared port)	Identifies the segment port of an EAPS ring that shares the common link.
Status (available with the detail keyword or by specifying a shared port)	<p>Up—Connectivity is established between the segment and the EAPS shared-port on the common link neighbor.</p> <p>Down—There is a break in the path between the segment and the EAPS shared-port on the common link neighbor.</p> <p>Blocking-Up—The path is Up, but due to the root blocker being in the Active state, this port is blocked to prevent a loop.</p> <p>Blocking-Down—The root blocker is in the Active state; however, the path is Down. Because the path is Down, there is no need to block the root blocker port to prevent a loop.</p> <p>[F]—The segment timer has expired but has not received an explicit link-down notification. The segment port remains in the Up state, with the timer expired flag set to True.</p>
EAPS Domain (available with the detail keyword or by specifying a shared port)	The EAPS domain assigned to the segment port.
Vlan-port count (available with the detail keyword or by specifying a shared port)	The total number of VLANs being protected on this segment port.
Adjacent Blocking Id (available with the detail keyword or by specifying a shared port)	<p>None—The neighbor on this port is not reporting a Controller in the Blocking state.</p> <p><i>Link-Id</i>—The neighbor on this port is a controller in the Blocking state with a link ID of <i>Link-Id</i>.</p>
Segment RB Id (available with the detail keyword or by specifying a shared port)	<p>None—The neighbor on this port is not aware of a root blocker in the network.</p> <p><i>RB-Id</i>—The neighbor on this port has determined that there is a root blocker in the network with a link ID of <i>RB-Id</i>.</p>
Vlan (available with the detail keyword or by specifying a shared port)	Displays a list of VLANs protected by the segment port.

Field	Description
Virtual-port Status (available with the detail keyword or by specifying a shared port)	This information appears for the Controller, when it is in either the Blocking or Preforwarding state. Active-Open—This VLAN or port is in the Forwarding state and has connectivity to the neighboring EAPS shared port via this port. Open—This VLAN or port is in the Forwarding state but does not have connectivity to the neighboring EAPS shared port via this port. Blocked—This VLAN or port is in the Blocking state to prevent a loop in the network. Down—This port's link is down. Active—At this moment, this VLAN or port is not being handled by EAPS shared port. Rather, this VLAN or port is being handled by the regular EAPS protocol.
Bvlan	When a common link connects an access VLAN (CVLAN or SVLAN) to a core VLAN (BVLAN), this field displays the BVLAN name. For more information, see the <i>Common Link Fault Detection and Response</i> section in the Switch Engine 32.3 User Guide .

Example

The following command displays shared-port information for all EAPS shared ports on a switch:

```
#show eaps shared-port
EAPS shared-port count: 1
-----
Link          Domain Vlan    RB      RB
Shared-port  Mode          Id  Up State    count  count Nbr State  Id
-----
10:1          Controller  1    Y  Ready      2       1    Yes None  None
Segment Timer expiry action: Send alert
-----
```

The following command displays detailed information for all EAPS shared ports:

```
#show eaps shared-port detail
EAPS shared-port count: 1
-----
Link          Domain Vlan    RB      RB
Shared-port  Mode          Id  Up State    count  count Nbr State  Id
-----
4:1          Controller  10    Y  Blocking    2       1    Yes Active  10
Segment Timer expiry action: Send alert
Segment Port: 5:7, Status: Blocking-Up
EAPS Domain:          d1
Vlan-port count:      1
Adjacent Blocking Id:  None
Segment RB Id:        None
Vlan                Virtual-port Status
p_1                  Blocked
Segment Port: 2:11,   Status: Down
EAPS Domain:          d2
Vlan-port count:      1
Adjacent Blocking Id:  20
Segment RB Id:        None
Vlan                Virtual-port Status
p_1                  Open
Vlan: p_1, Vlan-port count: 2, Active Open: None
Segment Port          Virtual-port Status
```

```
5:7          Blocked
2:11         Open
```

The following command displays detailed information for an EAPS shared port that is in the Blocking state:

```
# show eaps shared-port 1:24
-----
Link                               Domain Vlan      RB      RB      count  count Nbr State  Id
Shared-port  Mode                  Id  Up State      count  count Nbr State  Id
-----
1:24          Controller  10  Y  Blocking      3      5      Yes None  None
Segment Health Check interval:    1 sec
Segment Timeout:                  3 sec
Segment Fail Timer expiry action:  Send alert
Common Path Health Check interval: 1 sec
Common Path Timeout:              3 sec
Segment Port: 3:35  Status: Up
EAPS Domain:         d3
Vlan-port count:     3
Adjacent Blocking Id: None
Segment RB Id:        None
Segment Port: 3:36  Status: Up
EAPS Domain:         d2
Vlan-port count:     3
Adjacent Blocking Id: None
Segment RB Id:        None
Segment Port: 3:38  Status: Up
EAPS Domain:         d1
Vlan-port count:     5
Adjacent Blocking Id: None
Segment RB Id:        None
Vlan: data1,          Vlan-port count: 3,    Active Open: 3:38 Bvlan: metro1
Vlan: data2,          Vlan-port count: 3,    Active Open: 3:38 Bvlan: metro1
Vlan: data3,          Vlan-port count: 3,    Active Open: 3:38 Bvlan: metro2
Vlan: metro1,         Vlan-port count: 1,    Active Open: 3:38
Vlan: metro2,         Vlan-port count: 1,    Active Open: 3:38
-----
```



Note

The BVLAN information in the previous example appears only when a BVLAN configuration is present.

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms with the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and what licenses are appropriate for this feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

show eaps shared-port neighbor-info

```
show eaps shared-port {port} neighbor-info {detail}
```

Description

Displays shared-port information from neighboring shared links for one or more *EAPS* domains.

Syntax Description

<i>port</i>	Specifies a shared-port.
detail	Specifies to display the status of all segments and VLANs.

Default

N/A.

Usage Guidelines

If you enter the command without the detail keyword, the command displays a summary of status information for all configured EAPS shared ports from neighboring shared links. If you specify an EAPS shared-port, the command displays information about that specific port. Otherwise, the command displays information about all of the shared-ports configured on the switch.

You can use the detail keyword to display more detailed status information about the segments and VLANs associated with each shared port. For full details of the significant fields and values in the display output of the command, see the relevant tables in the [show eaps shared port {port} {detail}](#) command description.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show edp

```
show edp {ports [all | ports] {detail | vlan-id {mismatch {untagged}}  
         {neighbor nbr}}}
```

Description

Displays connectivity and configuration information for neighboring Extreme Networks switches.

Syntax Description

all	Specifies all ports.
<i>ports</i>	Specifies one or more ports or slots and ports.
vlan-id	Shows local and neighbor VLAN IDs.

mismatch	Shows only VLAN IDs that are mismatched to neighbor VLAN IDs.
untagged	Shows only mismatches for untagged ports.
neighbor	Shows only VLAN IDs for a specific neighbor.
<i>nbr</i>	Specifies the neighbor by Neighbor-ID (for example: 00:00:08:00:27:f3:f8:d9).
detail	Shows detailed information.

Default

N/A.

Usage Guidelines

port_list can be one or more port numbers. For a detailed explanation of port specification, see [Port Numbering](#) in [Command Reference Overview](#). To clear the counters, use the [clear lacp counters](#) command.

The neighbor-ID value is eight bytes. The first two bytes are always set to 00:00; the last six bytes are set to the neighbor's system MAC address.

Use the `show edp` command to display neighboring switches and configurations. This is most effective with Extreme Networks switches.

Example

The following example displays the configuration of the switch:

```
# show edp
EDP advert-interval      :60 seconds
EDP holddown-interval    :180 seconds
EDP enabled on ports     :1:1 1:2 1:3 1:4 1:5 1:6 3:1 3:2 3:3 3:4
```

The following example shows Extreme switch connectivity and configuration information for port 1:1:

```
# show edp ports 1:1
Port Neighbor
Port Vlans
=====
1:1 Oban 00:00:00:30:48:41:ed:97 1:1 54 1
=====
```

The following example displays detailed connectivity and configuration information of neighboring Extreme switches on port 1:1:

```
# show edp ports 1:1 detail
=====
Port 1:1: EDP is Enabled
Tx stats: sw-pdu-tx=2555      vlan-pdu-tx=1465      pdu-tx-err=0
Rx stats: sw-pdu-rx=2511      vlan-pdu-rx=2511      pdu-rx-err=0
Time of last transmit error: None
Time of last receive error: None
Remote-System: Oban          Age = 41
Remote-ID: 00:00:00:30:48:41:ed:97
```

```

Software version:      11.1.0.19
Remote-Port:          1:1
Port Type:            Ethernet
Auto Negotiation:      OFF
Flow Control:          SYMMETRIC/ASYMMETRIC
Duplex Speed:          Configured = HALF          Actual = HALF
Port Speed (MB):       Configured = ERROR         Actual = 100 Mbps
Remote-Vlans:
test (4094) Age = 41
=====

```

The following example shows local and neighbor VLAN IDs on port 3:

```

# show edp port 3 vlan-id
Local      Remote      |----- Local VLAN -----| |----- Remote VLAN
-----|
Port  Port  Neighbor ID      ID  Name              ID
Name              Age
-----|-----|-----|
3      1:3  00:00:08:00:27:88:2e:2b  200* v2-deepsouth      200* v2-
deepsouth              10
                        300 v3-articnorth      300 v3-
articnorth              10
-----|-----|-----|
3      1:3  00:00:08:00:27:f2:57:b5+      1
Default              12
                        200 v2-
deepsouth              300 v3-
articnorth
-----|-----|-----|
3      1:3  00:00:08:00:27:f3:e8:c9  200* v2-
deepsouth              300 v3-articnorth      300* v3-
articnorth              7
-----|-----|-----|

Flags: (*) Port is untagged in VLAN, (+) Neighbor does not send port tagged/untagged
information

```

The following example show mismatched neighbor VLAN IDs for untagged ports on port 3 :

```

# show edp port 3 vlan-id mismatch untagged
Local      Remote      |----- Local VLAN -----| |----- Remote VLAN
-----|
Port  Port  Neighbor ID      ID  Name              ID
Name              Age
-----|-----|-----|
3      1:3  00:00:08:00:27:f3:e8:c9  200* v2-
deepsouth              300 v3-articnorth      300* v3-
articnorth              11
-----|-----|-----|

```

History

This command was first available in ExtremeXOS 10.1.

Mismatched VLAN IDs check ability was added in ExtremeXOS 30.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show elrp

show elrp

Description

Displays ELRP information.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command displays the following:

- State of ELRP (enabled/disabled).
- Total number of ELRP sessions.
- ELRP packets transmitted.
- ELRP packets received.

In addition to the summary information at the top of the display, the `show elrp` command also displays the following information:

Client	Displays the name of the ELRP client.
VLAN	Displays the name of the VLAN with ELRP enabled.
Ports	Displays the set of VLAN ports used for packet transmission.
Interval	Displays the configured interval. An interval of 3 indicates that ELRP PDUs are transmitted every 3 seconds.
Count	Lists the configured number of ELRP PDUs that are transmitted. The PDUs are transmitted at the configured interval. This method of ELRP PDU transmission is used by <i>ESRP</i> in the pre-master state. A count of 0 indicates continuous PDU transmission. If the Cyclic value is Yes, the count is always 0.

Cyclic	Indicates whether ELRP PDUs are being continuously sent. The column shows Yes for the master VLAN because that VLAN is continuously sending ELRP PDUs for loop detection. When a VLAN is in the pre-master state, it only sends three ELRP PDUs before changing to master or slave. During this time the column shows No for that VLAN.
Pkts-Xmit	Displays the number of ELRP PDUs transmitted.
Pkts-Rcvd	Displays the number of ELRP PDUs received.
Action	Displays the configured action the switch takes when ELRP messages are received back indicating a detection of a network loop or no packets are received within the specified duration. The following list describes the actions: Print (P)—Specifies that the switch prints a message to the console.Log (L)—Specifies that the switch sends a message to the system log file.Trap (T)—Specifies that the switch sends a message to the <i>SNMP</i> manager.Callback (C)—Specifies a callback action. If you use ELRP with another protocol (for example ESRP), ELRP uses a callback action to notify the protocol of a loop detection.
Disable Port	Displays the configured hold time (number of seconds or permanent) for a port that was disabled with the <code>configure elrp-client periodic</code> command. When the time in seconds expires, the port is automatically enabled.

With hardware-assisted ELRP, one ACL egress counter is used to record ELRP PDUs transmitted per VLAN port, and the displayed Pkts-Xmit is calculated after all egress counters are queried from hardware. If show commands are executed very frequently, the statistics might not change every time.

Example

The following example displays summary ELRP status information on the switch:

```
# show elrp

ELRP Standalone Client:      Enabled
ELRP Inter-VLAN loop detection: On
ELRP client hardware assist: Enabled
Loopback port:              1

Number of ELRP sessions:    1
Number of ELRP pkts transmitted: 0
Number of ELRP pkts received: 0

Interval          Pkts    Pkts
Disable           Client VLAN Ports      (sec) Count Cyclic  Xmit    Rcvd Action Port
(sec)
-----
CLI   Default    All,          1.0    0 Yes    0       0 LT
-
                        All VXLAN RTEPs
-----
----
```

History

This command was first available in ExtremeXOS 11.1.

VXLAN remote endpoint information was added in ExtremeXOS 22.4.

Hardware-assisted ELRP information was added in ExtremeXOS 30.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show elrp disabled-ports

```
show elrp disabled-ports
```

Description

Displays information about ELRP disabled ports.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command displays the results of disabling ports using the `configure elrp-client periodic` command.

This command displays the following:

- Exclude EAPS ring ports—Whether EAPS ring ports can be excluded.
- Exclude VXLAN RTEPs—Whether VXLAN remote endpoints can be excluded.
- Exclude inter-VLAN loop ports—Whether inter-VLAN loop ports can be excluded.
- Excluded Ports—User defined ports that will not be disabled.
- Disabled Port—The port that ELRP disabled.
- Detected VLAN—The VLAN with looping ELRP PDU(s).
- Duration—The configured time to keep the port disabled.
- Time Disabled—The time when ELRP disabled the port.

Example

The following example shows ELRP information for disabled ports:

```
# show elrp disabled-ports

Exclude EAPS ring ports : No
Exclude VXLAN RTEPs    : Yes
Exclude inter-VLAN loop ports :No
```

Excluded Ports			

Disabled	Detected	Duration	Time
Disable			
Port/Virtual Port	Vlan	(sec)	Disabled
Direction			

--			

--			

History

This command was first available in ExtremeXOS 12.4.

The excluded port list was added in ExtremeXOS 12.5.3.

VXLAN remote endpoint exclusion information was added in ExtremeXOS 22.4.

Inter-VLAN loop ports exclusion information was added in ExtremeXOS 30.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show elrp dynamic-vlans

```
show elrp dynamic-vlans {vm-tracking | netlogin | mvrp | policy |
fabric-attach | all}
```

Description

This command shows ELRP configuration information for all, or specific types of, dynamic VLANs.

Syntax Description

dynamic-vlans	ELRP configuration options for dynamically created VLANs.
mvrp	Shows ELRP information for dynamic VLANs created by Multiple VLAN Registration Protocol (MVRP) only.
netlogin	Shows ELRP information for dynamic VLANs created by Network Login only.
vm-tracking	Shows ELRP information for dynamic VLANs created by virtual machine MAC tracking only.
policy	Specifies that the command applies to dynamic VLANs created by One Policy.

fabric-attach	Specifies that the command applies to dynamic VLANs created by Fabric Attach.
all	(Default) Shows ELRP information for all dynamic VLANs.

Default

If you do not specify the type of dynamic VLAN, information about all types appear.

Example

The following example shows ELRP information for all types of dynamic VLANs:

```
# show elrp dynamic-vlans

VM-Tracking :
ELRP on client ports : Disabled
ELRP on uplink ports : Disabled
ELRP on VXLAN RTEPs : Disabled
Transmit interval    : 1 second(s)
Loop detect action   : log and trap
Port disable on loop : Yes
Disabled port        : Ingress
Disabled duration    : 30 seconds

Netlogin :
ELRP on client ports : Disabled
ELRP on uplink ports : Disabled
ELRP on VXLAN RTEPs : Disabled
Transmit interval    : 1 second(s)
Loop detect action   : log and trap
Port disable on loop : Yes
Disabled port        : Ingress
Disabled duration    : 30 seconds

MVRP
Transmit interval    : 1 second(s)
Loop detect action   : log and trap
Port disable on loop : Yes
Disabled port        : Ingress
Disabled duration    : 30 seconds

Policy
Transmit interval    : 1 second(s)
Loop detect action   : log and trap
Port disable on loop : Yes
Disabled port        : Ingress
Disabled duration    : 30 seconds

Fabric-Attach
Transmit interval    : 1 second(s)
Loop detect action   : log and trap
Port disable on loop : Yes
Disabled port        : Ingress
Disabled duration    : 30 seconds
```

The following example shows ELRP information for dynamic VLANs created by virtual machine MAC tracking only:

```
# show elrp dynamic-vlans vm-tracking
VM-Tracking:
```

```
ELRP for dynamic VLANs      : Enabled
ELRP on client ports        : Enabled
ELRP on uplink ports        : Enabled
```

The following example shows ELRP information for dynamic VLANs created by One Policy:

```
# show elrp dynamic-vlans policy

Policy          : Enabled
Transmit interval : 1 second(s)
Loop detect action : log and trap
Port disable on loop : Yes
Disabled port    : Ingress
Disabled duration : 30 seconds
```

The following example shows ELRP information for dynamic VLANs created by Fabric Attach:

```
show elrp dynamic-vlans fabric-attach

Fabric-Attach      : Disabled
Transmit interval  : 1 second(s)
Loop detect action  : log and trap
Port disable on loop : Yes
Disabled port      : Ingress
Disabled duration   : 30 seconds
```

History

- This command was first available in ExtremeXOS 22.2.
- Loop detection action information was added in ExtremeXOS 22.3.
- The **policy** and **fabric-attach keywords** were added in ExtremeXOS 31.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show elsm ports

```
show elsm ports all | port_list
```

Description

Displays detailed information for one or more ELSM-enabled ports.

Syntax Description

all	Displays detailed ELSM information for all ports.
<i>port_list</i>	Displays detailed ELSM information for one or more ports.

Default

N/A.

Usage Guidelines

Use this command to display detailed information about the operational state of ELSM on the configured ports.

This command displays in a tabular format the following ELSM data for one or more ELSM-enabled ports on the switch:

- Port—The port number of the ELSM-enabled port.
- Link State—The state of the link between ELSM-enabled (peer) ports. The link state can be one of the following:
 - Ready—Indicates that the port is enabled but there is no link.
 - Active—Indicates that the port is enabled and the physical link is up.
- ELSM Link State—The current state of the ELSM logical link on the switch. The ELSM link state can be one of the following:
 - Up—Indicates that ELSM is enabled and the ELSM peer ports are up and communicating; the ELSM link state is up. In the up state, the ELSM-enabled port sends and receives hello messages from its peer.
 - Down—Indicates that ELSM is enabled, but the ELSM peers are not communicating; the ELSM link state is down. In the down state, ELSM transitions the peer port on this device to the down state. ELSM blocks all incoming and outgoing switching traffic and all control traffic except ELSM PDUs.
- ELSM State—The current state of ELSM on the port. The ELSM state can be one of the following:
 - Up—Indicates a healthy remote system and this port is receiving Hello+ messages from its peer.

If an ELSM-enabled port enters the Up state, the up timer begins. Each time the port receives a Hello+ message from its peer, the up timer restarts and the port remains in the Up state. The up timer is 6* hello timer, which by default is 6 seconds.

- Down—Indicates that the port is down, blocked, or has not received Hello+ messages from its peer.

If an ELSM-enabled port does not receive a hello message from its peer before the up timer expires, the port transitions to the Down state. When ELSM is down, data packets are neither forwarded nor transmitted out of that port.

- Down-Wait—Indicates a transitional state.

If the port enters the Down state and later receives a Hello+ message from its peer, the port enters the Down-Wait state. If the number of Hello+ messages received is greater than or equal to the hold threshold, the port transitions to the Up state. If the number of Hello+ messages received is less than the hold threshold, the port enters the Down state.

- Down-Stuck—Indicates that the port is down and requires user intervention.

If the port repeatedly flaps between the Up and Down states, the port enters the Down-Stuck state. Depending on your configuration, there are two ways for a port to transition out of this state:

By default, automatic restart is enabled, and the port automatically transitions out of this state. See the command `enable elsm ports auto-restart` for more information.

If you disabled automatic restart, and the port enters the Down-Stuck state, you can clear the stuck state and enter the Down state by using one of the following commands:

```
clear elsm ports port_list auto-restart
```

```
enable elsm ports port_list auto-restart
```

- Hello Transmit State—The current state of ELSM hello messages being transmitted. The transmit state can be one of the following:
 - HelloRx(+)—Specifies that the ELSM-enabled port is up and receiving Hello+ messages from its peer. The port remains in the HelloRx+ state and restarts the HelloRx timer each time it receives a Hello+ message. If the HelloRx timer expires, the hello transmit state enters HelloRX(-). The HelloRx timer is 6 * hello timer, which by default is 6 seconds.
 - HelloRx(-)—Specifies that the ELSM-enabled port either transitions from the initial ELSM state or is up and not receiving hello messages because there is a problem with the link or the peer is missing.
- Hello time—The current value of the hello timer, which by default is 1 second. The hello timer indicates the number of seconds between consecutive hello messages.
- Hold Threshold—The number of Hello+ messages required by the ELSM-enabled port to transition from the Down-Wait state to the Up state within the hold threshold.
- UpTimer Threshold—The number of hello times that span without receiving Hello+ packets before a port changes its ELSM state from Up to Down.
- Auto Restart—The current state of ELSM automatic restart on the port. The state of Auto Restart can be one of the following:
 - Enabled—If an ELSM-enabled port goes down, ELSM automatically brings up the down port. This is the default behavior.
 - Disabled If an ELSM-enabled port goes down, the port enters and remains in the Down-Stuck state until you clear the stuck state.

For more information about automatic restart, see the command `enable elsm ports auto-restart`.

- Sticky Threshold—Specifies the number of times a port can transition between the Up and Down states. The sticky threshold is not user-configurable and has a default value of 1. That means a port can transition only one time from the Up state to the Down state. If the port attempts a subsequent transition from the Up state to the Down state, the port enters the Down-Stuck state.
- Sticky Threshold Counter—The number of times the port transitions from the Up state to the Down state.
- Down Timeout—The actual waiting time (msecs or secs) before a port changes its ELSM state from Down to Up. When ELSM is enabled on a port and it is in a Down state, before it changes its ELSM state from Down to Up, it expects to receive at least a “Hold Threshold” number of Hello+ packets during the Down Timeout period after it receives the first Hello+ packet from its peer. It is equal to $[\text{Hello Time} * (\text{Hold Threshold} + 2)]$.
- Up Timeout—The actual waiting time (msecs or secs) before a port changes its ELSM state from Up to Down after receiving the last Hello+ packets. When a port is in an Up state, it expects to receive a Hello+ packet from its peer every “Hello Time” period to maintain its Up state. When it does not

receive a Hello+ packet after an “Up Timeout” period, it changes its ELSM state from Up to Down. It is equal to [Hello Time * UpTimer Threshold].

The remaining output displays counter information. Use the counter information to determine the health of the ELSM peers and how often ELSM has gone up or down. The counters are cumulative.

- RX Hello+—The number of Hello+ messages received by the port.
- Rx Hello- —The number of Hello- messages received by the port.
- Tx Hello+—The number of Hello+ messages sent by the port.
- Tx Hello- —The number of Hello- messages sent by the port.
- ELSM Up/Down Count—The number of times ELSM has been up or down.

To clear, reset the counters, use either the `clear elsm {ports port_list} counters` or the `clear counters` command.

Additional show Command

You can also use the `show ports {port_list} information {detail}` command to display ELSM information.

If you do not specify the detail parameter, the following columns display the current state of ELSM on the switch:

- Flags:
 - L—Indicates that ELSM is enabled on the switch.
 - - —Indicates that ELSM is disabled on the switch.
- ELSM:
 - up—Indicates that ELSM is enabled and the ELSM peer ports are up and communicating; the ELSM link state is up. In the up state, the ELSM-enabled port sends and receives hello messages from its peer.
 - dn—Indicates that ELSM is enabled, but the ELSM peers are not communicating; the ELSM link state is down. In the down state, ELSM transitions the peer port on this device to the down state. ELSM blocks all incoming and outgoing switching traffic and all control traffic except ELSM PDUs.
 - - —Indicates that ELSM is disabled on the switch.

If you specify the optional detail parameter, the following ELSM output is called out in written explanations versus displayed in a tabular format:

- ELSM Link State (displayed only if ELSM is enabled on the switch).
 - Up—Indicates that ELSM is enabled and the ELSM peer ports are up and communicating; the ELSM link state is up. In the up state, the ELSM-enabled port sends and receives hello messages from its peer.
 - Down—Indicates that ELSM is enabled, but the ELSM peers are not communicating; the ELSM link state is down. In the down state, ELSM transitions the peer port on this device to the down state. ELSM blocks all incoming and outgoing switching traffic and all control traffic except ELSM PDUs.
- ELSM:
 - Enabled—Indicates that ELSM is enabled on the switch.
 - Disabled—Indicates that ELSM is disabled on the switch.

Example

The following command displays detailed ELSM information for all configured ports on the switch:

```
show elsm ports all
```

The following is sample output from this command:

```
ELSM Info Port 4:4
Link State           : Active
ELSM Link State      : Up
ELSM State           : Up
Hello Transmit State  : HelloRx(+)
Hello Time           : 100 msec
Hold Threshold       : 2
UpTimer Threshold    : 6
Auto Restart         : Enabled
Down Timeout         : 400 msec
Up Timeout           : 600 msec
Rx Hello+            : 667960
Rx Hello-            : 0
Tx Hello+            : 667958
Tx Hello-            : 0
ELSM Up/Down Count   : UP: 0    DOWN: 0
ELSM Info Port 4:4
Link State           : Active
ELSM Link State      : Up
ELSM State           : Up
Hello Transmit State  : HelloRx(+)
Hello Time           : 100 msec
Hold Threshold       : 2
UpTimer Threshold    : 6
Auto Restart         : Disabled
Sticky Threshold     : 1
Sticky Threshold Counter : 0
Down Timeout         : 400 msec
Up Timeout           : 600 msec
Rx Hello+            : 708204
Rx Hello-            : 0
Tx Hello+            : 708201
Tx Hello-            : 0
ELSM Up/Down Count   : UP: 0    DOWN: 0
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show elsm

```
show elsm
```

Description

Displays summary information for all of the ELSM-enabled ports on the switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to display the operational state of ELSM on the configured ports.

If no ports are configured for ELSM, the switch does not display any information.

For ELSM-enabled ports, this command displays the following information in a tabular format:

- Port—The port number of the ELSM-enabled port.
- ELSM State—The current state of ELSM on the port. The ELSM state can be one of the following:
 - Up—Indicates a healthy remote system and this port is receiving Hello+ messages from its peer.

If an ELSM-enabled port enters the Up state, the up timer begins. Each time the port receives a Hello+ message from its peer, the up timer restarts and the port remains in the Up state.

- Down—Indicates that the port is down, blocked, or has not received Hello+ messages from its peer.

If an ELSM-enabled port does not receive a hello message from its peer before the up timer expires, the port transitions to the Down state. When ELSM is down, data packets are neither forwarded nor transmitted out of that port.

- Down-Wait—Indicates a transitional state.

If the port enters the Down state and later receives a Hello+ message from its peer, the port enters the Down-Wait state. If the number of Hello+ messages received is greater than or equal to the hello threshold (by default 2 messages), the port transitions to the Up state. If the number of Hello+ messages received is less than the hold threshold, the port enters the Down state.

- Down-Stuck—Indicates that the port is down and requires user intervention.

If the port repeatedly flaps between the Up and Down states, the port enters the Down-Stuck state. Depending on your configuration, there are two ways for a port to transition out of this state:

By default, automatic restart is enabled, and the port automatically transitions out of this state. See the command `enable elsm ports auto-restart` for more information.

If you disabled automatic restart, and the port enters the Down-Stuck state, you can clear the stuck state and enter the Down state by using one of the following commands:

```
clear elsm ports port_list auto-restart
```

```
enable elsm ports port_list auto-restart
```

- Hello time —The current value of the hello timer, which by default is 1 second. The hello timer indicates the number of seconds between consecutive hello messages.

Additional show Command

You can also use the `show ports {port_list} information {detail}` command to display ELSM information.

If you do not specify the detail parameter, the following columns display the current state of ELSM on the switch:

- Flags:
 - L—Indicates that ELSM is enabled on the switch.
 - - —Indicates that ELSM is disabled on the switch.
- ELSM:
 - up—Indicates that ELSM is enabled and the ELSM peer ports are up and communicating; the ELSM link state is up. In the up state, the ELSM-enabled port sends and receives hello messages from its peer.
 - dn—Indicates that ELSM is enabled, but the ELSM peers are not communicating; the ELSM link state is down. In the down state, ELSM transitions the peer port on this device to the down state. ELSM blocks all incoming and outgoing switching traffic and all control traffic except ELSM PDUs.
 - - —Indicates that ELSM is disabled on the switch.

If you specify the optional detail parameter, the following ELSM output is called out in written explanations versus displayed in a tabular format:

- ELSM Link State (displayed only if ELSM is enabled on the switch):
 - Up—Indicates that ELSM is enabled and the ELSM peer ports are up and communicating; the ELSM link state is up. In the up state, the ELSM-enabled port sends and receives hello messages from its peer.
 - Down—Indicates that ELSM is enabled, but the ELSM peers are not communicating; the ELSM link state is down. In the down state, ELSM transitions the peer port on this device to the down state. ELSM blocks all incoming and outgoing switching traffic and all control traffic except ELSM PDUs.
- ELSM:
 - Enabled—Indicates that ELSM is enabled on the switch.
 - Disabled—Indicates that ELSM is disabled on the switch.

Example

The following command displays summary configuration information for all of the ELSM-enabled ports on the switch:

```
show elsm
```

The following is sample output from this command:

```
Port      ELSM State  Hello Time
```

```
====
5:14 Up 1 (second)
5:18 Down 1 (second)
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show erps

show erps

Description

Display global information for *ERPS*.

Syntax Description

N/A.

Default

N/A.

Usage Guidelines

Use this command to display global information for ERPS.

Example

```
# show erps

ERPS Enabled: Yes
ERPS Display Config Warnings: On
ERPS Multicast Add Ring Ports: Off
ERPS Multicast Send IGMP Query: On
ERPS Multicast Temporary Flooding: Off
ERPS Multicast Temporary Flooding Duration: 15 sec
Number of ERPS instances: 1
# ERPS ring configuration :
```

Ring	State	Type	East	West	Control-Vlan	VID
R1	Protection	R r	21	+20	cv1	(1000)

```
where State: Init/Idle/Protection/Manual-Switch/Force-Switch/Pending
Type: (I) Interconnected node, (N) RPL Neighbor,
```

```
R) RPL Owner, (X) Ring node
Flags: (n) Non-revertive, (r) Revertive,
       (+) RPL Protection Port, (^) RPL Neighbor Port
       (f) Force Switch Port, (m) Manual Switch Port
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

show erps ring-name

```
show erps ring-name
```

Description

Display specific details for an *ERPS* ring.

Syntax Description

<i>ring-name</i>	Alphanumeric string that identifies the ERPS ring.
------------------	--

Default

N/A.

Usage Guidelines

Use this command to display specific details for an ERPS ring.

Example

The following example displays details for an ERPS ring named "R1":

```
# show erps "R1"

Name: R1
Operational State: Protection enabled      Node Type: RPL Owner,  Revertive
Configured State : Enabled

East Ring Port : 21  MepId: 1  Remote MepId: 3      Status: Blocked
West Ring Port : +20 MepId: 2  Remote MepId: 4      Status: Blocked

Periodic timer interval: 5000 millisec (Enabled)
Hold-off timer interval: 0    millisec (Enabled)
Guard timer interval   : 500  millisec (Enabled)
WTB timer interval     : 5500 millisec (Enabled)
WTR timer interval     : 1000 millisec (Enabled)
```



```
Ring MD Level      : 1
CCM Interval East  : 1000 millisec
CCM Interval West  : 1000 millisec
Notify Topology Change : -----
Subring Mode       : Virtual Channel

ERPS Control Vlan: cv1          VID:1000
Topology Change Propagation List: None
Topology Change Propagation : Disabled
ERPS Ring's Sub-Ring(s): None
ERPS Ring has following Protected Vlan(s):
  Vlan Name          VID
  pvl                1001
Number of Protected Vlans: 1
(+) RPL Protection Port, (^) RPL Neighbor Port
(f) Force Switch Port, (m) Manual Switch Port
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

show erps statistics

```
show erps ring-name statistics
```

Description

Display control packet and event statistics for an ERPS ring.

Syntax Description

<i>ring-name</i>	Alphanumeric string that identifies the ERPS ring.
------------------	--

Default

N/A.

Usage Guidelines

Use this command to display control packet and event statistics for an ERPS ring.

Example

The following example displays statistics for an ERPS ring named "R1":

```
# show erps "R1" statistics
```

port	Sent R-APS	Received R-APS	Dropped R-APS	Blocked events	Un-blocked events	SF	SF-clear
2:1	2309	3400	4	5	0	0	0
1:20	100	45	0	0	10	2000	100

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

show esrp

```
show esrp { {name} | {type [vpls-redundancy | standard]} }
```

Description

Displays *ESRP* configuration information for one or all ESRP domains on the switch.

Syntax Description

name	Specifies an ESRP domain name.
type	Specifies whether ESRP is standard or redundant VPLS.
vpls-redundancy	Specifies redundant VPLS.
standard	Specifies standard ESRP.

Default

Shows summary ESRP information.

Usage Guidelines

This command shows information about the state of an ESRP domain and its neighbors. This includes information about tracked devices.

In addition to ESRP information, ELRP status information is also displayed. This includes information about the master and pre-master states, number of transitions to the pre-master state, and the ports where ELRP is disabled.

The output varies depending upon the configuration and the state of the switch.

Example

The following command displays summary ESRP status information for the ESRP domains on the switch:

```
show esrp
```

The following is sample output from this command:

```
ESRP:                               Enabled
Configured Version:                 Extended
# ESRP domain configuration :
-----
Grp Ver VLAN      VID  DId  IP/IPX      State Master MAC Address Nbr
-----
ed2   0  E v2        2    2   2.2.2.3     Master 00:01:30:f9:9e:90 0
ed2   5  E v2        2    2   2.2.2.3     Aware  00:01:00:0D:9e:8a 0
ed3   0  E v3        3    3   0.0.0.0     Aware  00:01:00:0C:F0:D1 0
ed4   0  E v4        4    4   0.0.0.0     Slave  00:00:00:00:00:00 0
-----
# ESRP Port configuration:
-----
Port      Weight      Host      Restart
-----
6:1       0           H
6:2       10
6:3       0           R
```

The following command displays detailed ESRP status information for the specified ESRP domain on the switch (the election policy displayed is the default policy in extended mode):

```
show esrp ed2
```

The following is sample output from this command:

```
show esrp ed2
Domain:                ed2
Group:                 0
Operational Version:   extended
Vlan Interface:        v2
Vlan Tag:              2
Domain Id:             2
Rtif. Admin Status:    DOWN
Rtif. Virtual Mac :    00:e0:2b:00:00:80
IP Address:            2.2.2.3
Election Policy:
standby > sticky > ports > weight > track > priority > mac
-----
This System              Neighbor system
-----
State:                   Slave
Active Ports:            1
Tracked Active Vlan Ports: 0
Tracked IP Routes & Pings: 0
Priority:                255[255]
Sticky Flag:             0
MAC:                    00:01:30:f9:9e:90
Active Ports Weight:     10
Sequence Number:         8
Hand Shake Flag:         0
Restart Flag:            0
Timer Configuration:     Hello      2s (0)      Neighbor    6s (0)
```

```
PreMaster 4s(0) Neutral 4s(0)
NbrRestart 30s(0)
State Transition Counters: To Master 0 To Neutral 1
To PreMaster 0 To Aware 1
To Slave 1
Last State Change: Mon Apr 5 10:43:34 2004
ELRP in Premaster:      Enabled (Interval: 1, Count: 3)
ELRP in Master:         Enabled (Interval: 1)
Tracked Vlans: t_vlan
Tracked Pings: 40.0.1.2 / 3 second intervals / 5 misses / 10 successes
Tracked Ip Routes: 30.4.2.16/255.255.255.0
Tracked Environment: System Power : OK
Temperature : WARNING
-----
# Domain Member VLANs:
VLAN Name      VID      Virtual IP/IPX  State
-----
m_vlan1        1001     0.0.0.0         DOWN
m_vlan2        1002     0.0.0.0         DOWN
m_vlan3        1003     0.0.0.0         DOWN
```

History

This command was first available in ExtremeXOS 11.0.

ELRP status information was added in ExtremeXOS 11.1.

Ping success information added in ExtremeXOS 22.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show esrp aware

```
show esrp domain aware {selective-forward-ports | statistics}
```

Description

Displays all selective forwarding information for a given *ESRP*-aware domain.

Syntax Description

<i>domain</i>	Specifies the name of an ESRP domain.
selective-forward-ports	Specifies that the selective-forward-port table is the only table displayed.
statistics	Specifies that the selective-forward-port statistics table is the only table displayed.

Default

Disabled.

Usage Guidelines

An ESRP-aware switch floods ESRP PDUs from all ports in an ESRP-aware VLAN. This flooding creates unnecessary network traffic because some ports forward ESRP PDUs to switches that are not running the same ESRP groups. You can select the ports that are appropriate for forwarding ESRP PDUs by configuring selective forwarding on an ESRP-aware VLAN and thus reduce this excess traffic. Configuring selective forwarding creates a port list of only those ports that forward to the ESRP groups that are associated with an ESRP-aware VLAN. This aware port list is then used for forwarding ESRP PDUs.

The first of the two tables that this command displays shows Selective Forward Ports information:

- Group—The number of an ESRP group within the given domain.
- Port Count—The number of ports in the group that are selected to forward PDUs on the master VLAN.
- Selective Forward Ports—The list of ports in the group that are selected to forward PDUs on the master VLAN.

The second of the two tables displays statistical information about the activity of the ports:

- Group—The number of an ESRP group within the given domain.
- Master MAC—The MAC address for the master of the group.
- Rx—The number of PDUs received matching the domain/group pair.
- Fwd—The number of PDUs received and forwarded matching the domain/group pair.
- FDB Flush—The number of FDB Flush requests received from the ESRP Master for this domain/group pair.
- Fwd Ports—Selective or Default.

Selective describes the group as having a configured aware port list for selective forwarding of PDUs on the Master VLAN. The list of ports is displayed in the first table above.

Default describes the group as one where all the ports on the master VLAN forward the ESRP PDUs that are received for the domain/group pair. Because there is no selective forwarding configuration for this group, there is no entry in the first table.

Example

The following command displays the ESRP aware information for the domain d1.

```
show esrp d1 aware
Domain:          d1
Vlan:            vesrp1
```

Group	Port Count	Selective Forward Ports
0	5	5:1, 5:2, 7:31, 7:32: 8:1
3	2	5:1, 8:1

Group	Master MAC	Fwd Ports	Rx	Fwd
0	00:12:00:33:44:55		10	
10	1	selective		
1	00:22:00:12:21:1F		77	

```

77                               3
default
3          00:02:00:13:11:11          99
99                               3          selective

```

History

This command was first available in Extreme XOS 12.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show esrp counters

```
show esrp {name} counters
```

Description

Displays *ESRP* counter information for ESRP or for a specified ESRP domain.

Syntax Description

<i>name</i>	Specifies an ESRP domain name.
-------------	--------------------------------

Default

Displays summary ESRP counter information.

Usage Guidelines

The `show esrp counters` command displays information about the number of:

- Failed received protocol packets
- Failed sent protocol packets
- Dropped protocol packets belonging to unknown ESRP domains
- Dropped protocol packets due to invalid Extreme Encapsulation Protocol (EEP) data
- Dropped packets due to old sequence numbers
- Dropped packets due to an invalid 802.1Q tag
- Dropped packets because the packet length was truncated (packet length is less than expected)
- Dropped packets due to failed checksum verification

The `show esrp {name} counters` command displays information about the number of times ESRP, ESRP-aware, and ESRP error packets were transmitted and received.

Example

The following command displays ESRP counter information:

```
show esrp counters
```

The following is sample output from this command:

```
Current-time:                Sun Nov 16 00:25:08 2003
esrpStatsRxHelloFailed      = 0
esrpStatsTxHelloFailed      = 0
esrpStatsUnknownDomain      = 0
esrpStatsUnsupportedEEPVersion = 0
esrpStatsInvalidEPLength    = 0
esrpStatsNotInTimeWindow    = 0
esrpStatsInvalid8021Qtag     = 0
esrpStatsInvalidSNAPType     = 0
esrpStatsUndersizePkt        = 0
esrpStatsInvalidChecksum     = 0
esrpStatsWrongDigest         = 0
```

The following command displays counter information for ESRP domain ed5:

```
show esrp ed5 counters
```

The following is sample output from this command:

```
Domain:  ed5                Current-time:  Sun Nov 16 00:25:27 2003
Rx-Esrp-Pkts      = 628      Tx-Esrp-Pkts      = 630
Rx-Aware-Esrp-Pkts = 112      Tx-Aware-Esrp-Pkts = 34
Rx-Err-Pkts       = 0 Tx-Err-Pkts       = 0
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show ethernet oam

```
show ethernet oam {ports [port_list] {detail}}
```

Description

Displays Ethernet OAM information.

Syntax Description

<i>port_list</i>	Specifies the particular ports.
detail	Specifies that detailed information be displayed.

Default

N/A.

Usage Guidelines

Use this command to display basic Ethernet OAM information for specified ports on the switch. If you do not specify the port(s), information for all ports is displayed.

Use the detail option for additional information.

When operating as a stack master, the switch can process this command for ports on supported platforms.

Example

The following command displays basic Ethernet OAM information for all ports:

```
show ethernet oam
```

Following is sample output from the command:

```
switch # show ethernet oam
=====
Port  Flags   Tx Cnt Rx Cnt Tx Err Rx Err
=====
1     E--u    2      2      0      0
2     ---u    0      0      0      0
3     E-Ru    2      2      0      0
4     ---u    0      0      0      0
5     EU-u    0      0      0      0
6     ---u    0      0      0      0
7     ---u    0      0      0      0
8     ---u    0      0      0      0
9     ---u    0      0      0      0
10    ---u    0      0      0      0
11    ---u    0      0      0      0
12    ---u    0      0      0      0
13    ---u    0      0      0      0
14    ---u    0      0      0      0
15    ---u    0      0      0      0
16    ---u    0      0      0      0
17    ---u    0      0      0      0
18    ---u    0      0      0      0
19    ---u    0      0      0      0
20    ---u    0      0      0      0
21    ----    0      0      0      0
22    ----    0      0      0      0
23    ----    0      0      0      0
24    ----    0      0      0      0
25    ----    0      0      0      0
26    ----    0      0      0      0
-----
Flags   : (E) OAM Enabled, (U) OAM Operationally Up,
          (R) Remote Port Fault Exists,
          (u) Unidirectional OAM Supported
```


The following command displays detailed information for port 1:

```
show ethernet oam port 1 detail
```

Following is sample output from the command:

```
switch # show ethernet oam port 1 detail
Port Number      : 1
Admin Status     : Enabled      Unidirectional OAM : Supported
Oper Status      : Disabled     Remote Fault       : Not Exists
Tx Pkts          : 2527         Rx Pkts            : 2550
Tx Error         : 0            Rx Error           : 0
```

History

This command was first available in ExtremeXOS 12.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show fabric attach agent

```
show fabric attach agent
```

Description

Displays known fabric attach agent information.

Syntax Description

agent	Displays Fabric Attach agent information.
--------------	---

Default

N/A.

Example

The following example shows Fabric Attach agent information:

```
# show fabric attach agent
Fabric Attach Agent Status
-----
Service Status:      Enabled
Element Type:        Disabled
Zero Touch Status:    Not Supported
Auto Provision Setting: Not Supported
Provision Mode:       Not Supported
Client Proxy Status:  Enabled
Standalone Proxy Uplink: 10
Agent Timeout:       240 seconds
```

```

Extended Logging Status: Not Supported
Primary Server ID:      00-04-96-a0-a4-93-00-01-00-05
Server Descr: ExtremeXOS (X670G2-48x-4q) version 22.5.0.43 xos_22.5
                    by cmbuild on Wed Mar 28 08:57:47 EDT 2018
Forward Management VLAN: Enabled

```

History

This command was first available in ExtremeXOS 22.5.

The information about whether or not Fabric Attach proxy switches send Management VLAN data to clients was added in ExtremeXOS 30.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show fabric attach assignments

```

show fabric attach [{vlan} | {vlan} vlan_name | vlan vlan_id]
                    assignments

```

Description

Displays known fabric attach assignment information.

Syntax Description

assignments	Display Fabric Attach assignments.
vlan	Designates specifying the VLAN.
<i>vlan_name</i>	Specifies VLAN name.
<i>vlan_id</i>	Specifies VLAN ID.

Default

N/A.

Example

The following example shows Fabric Attach assignments information:

```

# show fabric attach assignments
Fabric Attach Mode: Proxy
Port    VLAN    VLAN Name    Type    ISID/NSI    Status
-----
          100    VLAN_0100    Static    100         Active
          400    v4           Static    400         Active
          1000   VLAN_1000    Static    1000        Active
5        400    v4           Dynamic    404040      Rejected
9         20    SYS_VLAN_0020 Dynamic    20          Active

```

9	2500	SYS_VLAN_2500	Dynamic	25	Active
9	3500	SYS_VLAN_3500	Dynamic	35	Active

History

This command was first available in ExtremeXOS 30.2.

VLAN syntax was clarified in ExtremeXOS 30.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show fabric attach elements

show fabric attach **elements**

Description

Displays known fabric attach elements information.

Syntax Description

elements	Display Fabric Attach elements (neighbors).
-----------------	---

Default

N/A.

Example

The following example shows Fabric Attach element information:

```
# show fabric attach elements
Fabric Attach Mode: Standalone Proxy
```

System Id	Port	Type	Mgmt VLAN	Tag	Auto Provision
00-21-70-f2-37-05-00-00-00-00	2	Virtual Switch	1212	Mix	Disabled
6c-a8-49-5a-44-00-00-00-00-d3	20	Server (No Auth)	1212	Mix	Disabled

History

This command was first available in ExtremeXOS 22.4.

The the term **neighbors** was changed to **elements** in ExtremeXOS 22.5.

Standalone proxy mode information was added in ExtremeXOS 30.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show fabric attach port

```
show fabric attach ports [port_list | all]
```

Description

Displays the Fabric Attach port configuration.

Syntax Description

ports	Specifies the port to configure.
<i>port_list</i>	Specifies a list of ports to configure.
all	Specifies configuring all ports in the system.

Default

If no mapping is present, no information will be displayed for the client type.

Usage Guidelines

Use this command to display configured Fabric Attach ports.

Example

The following example displays configured Fabric Attach ports:

```
# show fabric attach ports 25
Port      Disabled Managment Forward Authentication
-----
25        Yes      Off          Disabled
```

History

This command was first available in ExtremeXOS 31.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show fabric attach ports authentication

```
show fabric attach ports [port_list | all] authentication {detail}
```

Description

Shows Fabric Attach authentication information.

Syntax Description

ports	Specifies ports to show.
<i>port_list</i>	Specifies the list of ports to show.
all	Shows information for all ports in the system.
authentication	Specifies showing Fabric Attach authentication information.
detail	Specifies showing detailed Fabric Attach authentication information.

Default

N/A.

Usage Guidelines

This command shows the configuration settings from the `configure fabric attach ports [port_list | all] authentication [disable | enable | key {key | default | encrypted encrypted_key}]` command.

Example

The following example shows Fabric Attach authentication information on port 1:

```
# show fabric attach ports 1 authentication
Port      Authentication
-----
1         Enabled
```

The following example shows detailed Fabric Attach authentication information on all ports:

```
show fabric attach ports all authentication detail
Port      Authentication  Encrypted Key
-----
1         Enabled      "#$GuZJvU40xnRwKd+PUlsBxSwfpQYr5w=="
2         Disabled
"#4MoonVveq8vrk9BrKO9I673CCDfnWUEcAcKeTZcvQup3XFApn36GWvMGXA13TpufrUXPZPV2geztybwV/
0xfLoFt+zg/4g=="
```

History

This command was first available in ExtremeXOS 30.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show fabric attach statistics

show fabric attach **statistics**

Description

Displays known fabric attach statistics information.

Syntax Description

statistics	Displays Fabric Attach statistics.
-------------------	------------------------------------

Default

N/A.

Usage Guidelines

To clear these statistics, use the `clear counters` command.

Example

The following example shows Fabric Attach statistics information:

# show fabric attach statistics				
Port	Element Received	Element Timed Out	Element Deleted	Auth Failed
-----	-----	-----	-----	-----
1:20	4000000000	4000000000	4000000000	N/A
1:2	25	0	3	N/A
2:32	1234	0	12	N/A

History

This command was first available in ExtremeXOS 22.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show fabric attach zero-touch-client

show fabric attach **zero-touch-client**

Description

Displays the configured ZTC mappings for Fabric Attach Zero Touch Clients.

Syntax Description

zero-touch-client	Displays Zero Touch Clients.
--------------------------	------------------------------

Default

If no mapping is present, no information will be displayed for the client type.

Usage Guidelines

This command is only used when mappings are present.

Example

The following example displays configured Fabric Attach Zero Touch Clients:

```
# show fabric attach zero-touch-client
```

Client Name	VLAN Name	NSI/I-SID	Priority	Status
wap-type1	VLAN_123	11111	dot1p	Enabled
camera	VLAN_200	2000	5	Disabled
ona-spb-over-ip	VLAN_4001	40001	7	Enabled

History

This command was first available in ExtremeXOS 31.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show failsafe-account

```
show failsafe-account
```

Description

Displays whether the user configured a username and password for the failsafe account or shows the configured connection type access restrictions.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to view the failsafe account configuration.

The command shows the access permissions and whether or not the user configured a username and password. It does not show the configured username or password.

Example

The following command displays the failsafe account configuration:

```
show failsafe-account
```

Output from this command looks similar to the following when a failsafe account username and password have been configured with all connections types permitted for failsafe account access:

```
The Failsafe Account is enabled for these connection types:
- Serial Console
- Control Fabric (inter-node)
```

Output from this command looks similar to the following when the factory default (i.e. no user-specified username and password are configured) is in effect:

```
The Failsafe Account is disabled.
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show fans

```
show fans {detail}
```

Description

Displays the status of the fans in the system.

Syntax Description

detail	The detail option is reserved for future use.
---------------	---

Default

N/A.

Usage Guidelines

Use this command to view detailed information about the health of the fans.

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

The switch collects and displays the following fan information:

- State—The current state of the fan. Options are:
 - Empty: There is no fan installed.
 - Failed: The fan failed.
 - Operational: The fan is installed and working normally.
- NumFan—The number of fans in the fan tray.
- Fan Name, displayed as Fan-1, Fan-2, and so on—Specifies the individual state for each fan in a fan tray and its current speed in revolutions per minute (rpm).



Note

For the ExtremeSwitching X435-24P-4S switches, you cannot view the fan speed.

Bridge Port Extenders (BPEs)

Fan information for bridge port extenders (BPEs) does not appear in this command. Use the `show vpx bpe {slot slot_num} {environment}` command instead.

Example

The following command displays the status of the installed fans. If a fan is not installed, the state of the fan is Empty.

```
show fans
```

The following is sample output from an ExtremeSwitching switch:

```
FanTray information:
State:                Operational
NumFan:               6
PartInfo:             0931G-00064 450237-00-05
Revision:             5.0
Fan-1:               Operational at 14894 RPM
Fan-2:               Operational at 15360 RPM
Fan-3:               Operational at 15360 RPM
Fan-4:               Operational at 9637 RPM
Fan-5:               Operational at 9637 RPM
Fan-6:               Operational at 9637 RPM
```

The following is a sample output from a SummitStack:

```
FanTray-1 information:
State:                Operational
NumFan:               1
Fan-1:               Operational at 1000 RPM
FanTray-2 information:
```

```

State:                Operational
NumFan:               1
Fan-1:               Operational at 1000 RPM
FanTray-3 information:
State:                Operational
NumFan:               1
Fan-1:               Operational at 1000 RPM
FanTray-4 information:
State:                Operational
NumFan:               1
Fan-1:               Operational at 1000 RPM
FanTray-5 information:
State:                Empty
FanTray-6 information:
State:                Empty
FanTray-7 information:
State:                Empty
FanTray-8 information:
State:                Empty

```

History

This command was first available in an ExtremeXOS 10.1.

Information about the current speed in RPM was added to the `show fans` output in ExtremeXOS 11.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show fdb

```

show fdb {blackhole {netlogin [all | mac-based-vlans]} | netlogin [all
| mac-based-vlans] | permanent {netlogin [all | mac-based-vlans]}
| mac_addr {netlogin [all | mac-based-vlans]} | ports port_list
{netlogin [all | mac-based-vlans]} | [ {vlan} vlan_name | vlan
vlan_list] {netlogin [all | mac-based-vlans]} | [{vpls} {vpls_name} |
vxlan { vni } | virtual-network vn_name}}

```

Description

Displays *FDB* entries.

Syntax Description

blackhole	Displays the blackhole entries. (All packets addressed to these entries are dropped.)
slot	Specifies a slot in the switch.
num_entries	Specifies the maximum number of hardware entries to display. The range is 1 to 25.
netlogin all	Displays all FDBs created as a result of the netlogin process.

netlogin mac-based-vlans	Displays all netlogin MAC-based VLAN FDB entries.
permanent	Displays all permanent entries, including the ingress and egress QoS profiles.
<i>mac_addr</i>	Specifies a MAC address, using colon-separated bytes, for which FDB entries should be displayed.
<i>port_list</i>	Displays the entries for one or more ports or ports and slots.
<i>vlan_name</i>	Displays the entries for a specific VLAN.
<i>vlan_list</i>	Displays a VLAN list of IDs.
<i>vpls_name</i>	Specifies a specific VPLS for which to display entries.
vxlan	MAC addresses reachable through VXLAN-encapsulated Tunnels.
<i>vni</i>	MAC Addresses reachable through VXLAN tunnels encapsulated with the specified Virtual Network Identifier value.
virtual-network	MAC addresses associated with a Virtual Overlay Network learning domain.
<i>vn_name</i>	Alphanumeric string identifying the Virtual Network whose associated MAC addresses need to be displayed.

Default

All.

Usage Guidelines

The pulling of MAC addresses for display purposes is given a lower priority to the actual data path learning. Eventually all the MAC addresses are learned in a quiescent system.

The show fdb command output displays the following information:

Mac	The MAC address that defines the entry.
Vlan	The PVLAN or VLAN for the entry.
Age	The age of the entry, in seconds (does not appear if the keyword permanent is specified).

Flags	Flags that define the type of entry: <ul style="list-style-type: none"> • b - Ingress Blackhole • B - Egress Blackhole • D - Drop entry for an isolated subscriber VLAN • d - Dynamic • h - hardware aging • i - an entry also exists in the IP FDB • l - lockdown MAC • L - lockdown-timeout MAC • m - MAC • M - Mirror • n - <i>NetLogin</i> • o - IEEE 802.1ah backbone MAC • P - PVLAN created entry • p - Permanent • s - Static • v - NetLogin MAC-Based VLAN • x - an entry also exists in the IPX FDBs. • X - VXLAN
Port List	The ports on which the MAC address has been learned.

Example

The following example shows how the FDB entries appear for all options except the hardware option:

```
# show fdb
Mac                Vlan      Age  Flags      Port / Virtual Port List
-----
00:0c:29:4b:34:cf  v101(0101) 0041 d m      D 1:2
00:0c:29:4b:34:cf  v100(0100) 0041 d m      P 1:2
00:0c:29:d2:2d:48  v102(0102) 0045 d m      1:3
00:0c:29:d2:2d:48  v100(0100) 0045 d m      P 1:3
00:0c:29:f1:f2:f5  v100(0100) 0045 d m      1:1
00:0c:29:f1:f2:f5  v102(0102) 0045 d m      P 1:1
00:0c:29:f1:f2:f5  v101(0101) 0045 d m      P 1:1
Flags : d - Dynamic, s - Static, p - Permanent, n - NetLogin, m - MAC, i - IP,
x - IPX, l - lockdown MAC, L - lockdown-timeout MAC, M- Mirror, B - Egress
        Blackhole,
b - Ingress Blackhole, v - MAC-Based VLAN, P - Private VLAN, T - VLAN
        translation,
D - drop packet, h - Hardware Aging, o - IEEE 802.1ah Backbone MAC.
S - Software Controlled Deletion, r - MSRP
X - VXLAN, E - EVPN
Total: 3 Static: 0 Perm: 0 Dyn: 3 Dropped: 0 Locked: 0 Locked with Timeout: 0
FDB Aging time: 300
FDB VPLS Aging time: 300
```

The following example output shows where the port tag is displayed in parentheses:

```
# show fdb
Mac                Vlan      Age  Flags      Port / Virtual Port List
-----
00:00:00:00:04:0a  test(0200) 0057 d m      3(0010)
00:00:00:00:04:0b  test(0200) 0300 d m      3(0011)
00:01:02:03:04:05  test(0200) 0000 spm      3(0010)
Flags : d - Dynamic, s - Static, p - Permanent, n - NetLogin, m - MAC, i - IP,
```

```

    x - IPX, l - lockdown MAC, L - lockdown-timeout MAC, M- Mirror, B - Egress
    Blackhole,
    b - Ingress Blackhole, v - MAC-Based VLAN, P - Private VLAN, T - VLAN translation,
    D - drop packet, h - Hardware Aging, o - IEEE 802.1ah Backbone MAC,
    S - Software Controlled Deletion, r - MSRP,
    X - VXLAN, E - EVPN
    Total: 3 Static: 0 Perm: 0 Dyn: 3 Dropped: 0 Locked: 0 Locked with Timeout: 0
    FDB Aging time: 300
    FDB VPLS Aging time: 300

```

The following show fdb example includes VXLAN output:

```

# show fdb vxlan
Mac                Vlan Age Flags Port / Virtual Port List
-----
00:00:00:00:00:21 vlan101 0000 s mX VR-Default:20.20.20.1
01:00:5E:00:00:01 vlan101 0000 s mX VR-Default:20.20.20.1
                                s mX VR-Default:30.30.30.1
FF:FF:FF:FF:FF:FF vlan101 0000 s mX VR-Default:20.20.20.1
                                s mX VR-Default:30.30.30.1
                                s mX VR-Default:40.40.40.1
Unknown-Unicast vlan101 0000 s mX VR-Default:20.20.20.1
                                s mX VR-Default:30.30.30.1
                                s mX VR-Default:40.40.40.1
00:00:00:00:00:22 vlan102 0000 s mX VR-Default:20.20.20.1
00:00:00:00:00:23 vlan103 0000 s mX VR-Default:20.20.20.1
00:00:00:00:00:24 vlan104 0000 s mX VR-Default:20.20.20.1
Flags : d - Dynamic, s - Static, p - Permanent, n - NetLogin, m - MAC, i - IP,
x - IPX, l - lockdown MAC, L - lockdown-timeout MAC, M- Mirror, B - Egress Blackhole,
b - Ingress Blackhole, v - MAC-Based VLAN, P - Private VLAN, T - VLAN translation,
D - drop packet, h - Hardware Aging, o - IEEE 802.1ah Backbone MAC,
S - Software Controlled Deletion, r - MSRP, X - VXLAN, E - EVPN
Total: 7 Static: 7 Perm: 0 Dyn: 0 Dropped: 1 Locked: 0 Locked with Timeout: 0
FDB Aging time: 300

```

History

This command was first available in ExtremeXOS 10.1.

The **stats** and **netlogin** parameters were first available in ExtremeXOS 11.3.

The blackhole output under the **b** and **B** flags was first available for all platforms in ExtremeXOS 12.1.

The **o** flag was first available in ExtremeXOS 12.4.

The `vlan_list` variable was added in ExtremeXOS 16.1.

The command was enhanced to show the MAC to IP binding entries in ExtremeXOS 21.1. In addition to retrieving all FDB entries known to the switch, you can query the filtering database for all fdb entries learned within a virtual network (key: <vn_name>), all entries learnt on VXLAN tunnels (token: vxlan) and entries learned on VXLAN tunnels on a specific VNI (key:VNI).

Flag for EVPN-learned entries was added in ExtremeXOS 22.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show fdb mac-tracking configuration

```
show fdb mac-tracking configuration
```

Description

Displays configuration information for the MAC address tracking feature.

Syntax Description

This command has no arguments or variables.

Default

The MAC address tracking table is empty.

Usage Guidelines

None.

Example

The following example displays the contents of the MAC address tracking table:

```
# show fdb mac-tracking configuration
MAC-Tracking enabled ports: 1-3,10,20
SNMP trap notification      : Enabled
MAC address tracking table (4 entries):
00:30:48:72:ee:88
00:21:9b:0e:ca:32
00:12:48:82:9c:56
00:30:48:84:d4:16
```

History

This command was first available in ExtremeXOS 12.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show fdb mac-tracking statistics

```
show fdb mac-tracking statistics {mac_addr} {no-refresh | refresh}
```

Description

Displays statistics for the MAC addresses that are being tracked.

Syntax Description

<i>mac_addr</i>	Specifies a MAC address, using colon-separated bytes, for which <i>FDB</i> entries should be displayed.
no-refresh	Specifies a static snapshot of data instead of the default dynamic display.
refresh	Specifies a continuous refresh of output.

Default

N/A.

Usage Guidelines

Use the keys listed below the display to clear the statistics counters or page up or down through the table entries.

Example

The following example displays statistics for the entries in the MAC address tracking table:

```
# show fdb mac-tracking statistics
MAC Tracking Statistics      Fri Mar 20 15:25:01 2009
Add      Move      Delete
MAC Address      events      events      events
=====
00:00:00:00:00:01      0          0          0
00:00:00:00:00:02      0          0          0
00:00:00:00:00:03      0          0          0
00:00:00:00:00:04      0          0          0
00:00:00:00:00:05      0          0          0
00:00:00:00:00:06      0          0          0
00:00:00:00:00:07      0          0          0
00:00:00:00:00:08      0          0          0
00:00:00:00:00:09      0          0          0
00:00:00:00:00:10      0          0          0
00:00:00:00:00:11      0          0          0
00:00:00:00:00:12      0          0          0
00:00:00:00:00:13      0          0          0
00:00:00:00:00:14      0          0          0
00:00:00:00:00:15      0          0          0
00:00:00:00:00:16      0          0          0
00:00:00:00:00:17      0          0          0
00:00:00:00:00:18      0          0          0
=====
0->Clear Counters  U->page up  D->page down  ESC->exit
```

History

This command was first available in ExtremeXOS 12.3.

The **refresh** keyword was added in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show fdb static-mac-move configuration

```
show fdb static-mac-move configuration
```

Description

Displays the configuration for the feature that reports the discovery of MAC addresses that are duplicates of statically configured MAC addresses.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following example shows the command display:

```
# show fdb static-mac-movement configuration
Static MAC Movement Notification: Enabled
MAC learning Packets Count       : 5
```

History

This command was first available in ExtremeXOS 12.7.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show fdb stats

```
show fdb stats [{ports {all | port_list} | vlan {all} | [ {vlan}
                vlan_name | vlan vlan_list] {no-refresh | refresh}}]
```

Description

Displays *FDB* entry statistics for the specified ports or VLANs in either a dynamic or a static report.

Syntax Description

all	Requests statistics for all ports or all VLANs.
<i>port_list</i>	Specifies which ports are to be included in the statistics display.
<i>vlan_name</i>	Specifies a single <u>VLAN</u> to be included in the statistics display.
<i>vlan_list</i>	Specifies a VLAN list of IDs to be included in the statistics display.
refresh	Specifies a continuous refresh of output.

Default

Summary FDB statistics for the switch.

Usage Guidelines

The dynamic display remains visible and continues to update until you press **[Esc]**.

The `show fdb stats` command output displays the following information:

Port	When you chose to display statistics for ports, this column displays port numbers.
Link State	When you chose to display statistics for ports, this column displays the link states, which are described at the bottom of the display.
VLAN	When you chose to display statistics for VLANs, this column displays VLAN names.
MAC Addresses	This column displays the total number of MAC addresses for each port or VLAN.
Dynamic	This column displays the total number of MAC addresses that were learned dynamically for each port or VLAN.
Static	This column displays the total number of MAC addresses that are configured on this switch for each port or VLAN.
Dropped	This column displays the total number of dynamic MAC addresses that were discovered, but not stored in the FDB. Discovered MAC addresses might be dropped because a configured learning limit is reached, the FDB is in lockdown, or a port forwarding state is in transition. Some conditions that lead to dropped MAC addresses can produce log messages or <u>SNMP</u> traps.

Example

The following command example displays summary FDB statistics for the switch:

```
torino1.1 # show fdb stats
Total: 4 Static: 3 Perm: 3 Dyn: 1 Dropped: 0
FDB Aging time: 300
FDB VPLS Aging time: 300
torino1.2 #
```

The following command example displays FDB statistics for ports 1 to 16 on slot 1:

```
# show fdb stats ports 1:1-1:16
FDB Stats                               Mon Mar 15 15:30:49 2010
Port      Link  MAC
State    Addresses      Dynamic      Static      Dropped
=====
1:1       A      2394        2389        5          2
1:2       A      37          37          0          0
1:3       A      122         121         1          452
1:4       R      0           0           0          0
1:5       R      0           0           0          0
1:6       A      43          43          0          0
1:7       A      118         118         0          0
1:8       R      0           0           0          0
1:9       R      0           0           0          0
1:10      A      8           8           0          0
1:11      A      2998        2990        8          1
1:12      A      486         486         0          0
1:13      R      0           0           0          0
1:14      A      42          42          0          0
1:15      A      795         795         0          0
1:16      A      23          23          0          2
=====
Link State: A-Active, R-Ready, NP-Port Not Present, L-Loopback
U->page up  D->page down  ESC->exit
```

The following command example displays FDB statistics for all VLANs:

```
# show fdb stats vlan all
FDB Stats                               Mon Mar 15 15:30:49 2010
VLAN      MAC Addresses      Dynamic      Static      Dropped
=====
SV_PPPOE      2394        2389        5          2
NV_PPPOE      122         121         1          452
=====
U->page up  D->page down  ESC->exit
```

History

The dynamic display for this command was first available in ExtremeXOS 12.4.2.

The **no-refresh** keyword was removed in ExtremeXOS 16.1.

The `vlan_list` variable was added in ExtremeXOS 16.1.

The **refresh** keyword was added in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show flow-redirect

```
show flow-redirect {flow_redirect_name}
```

Description

Displays nexthop ipaddresses, up/down status, health-checking (ping/ARP/ND) and [ACL](#) bindings.

Syntax Description

<code>flow_redirect_name</code>	Specifies the name of the flow redirection policy.
---------------------------------	--

Default

N/A.

Usage Guidelines

None.

Example

The following example displays information for all redirection policies:

```
Switch.1 # show flow-redirect
Name          Nexthop  Active          VR Name      Inactive  Health
Count        IP Address      Nexthops    Check
=====
pkh           2          2001:400::100  VR-Default  Forward  PING
ND: Neighbor Discovery
```

The next example displays an IPv6 redirection policy for a longer IPv6 address, which causes a two-line display for the related redirection policy:

```
Switch.13 # sh flow-redirect
Name          Nexthop  Active          VR Name      Inactive  Health
Count        IP address      Nexthops    Check
=====
pbr1          2          2004:1000:1000:1000::10
VR-Default   Forward    PING
ND: Neighbor Discovery
```

This example displays information for a specified IPv6 redirection policy:

```
* Switch.14 # show flow-redirect "pbr1"
Name          : pbr1          VR Name      : VR-Default
Inactive Nexthops: Forward    Health Check  : PING
Nexthop Count  : 2
Active IP Address : 2004:1000:1000:1000::10
Index   State   Priority  IP Address      Status Interval Miss Success
=====
0       Disabled 200      2003::10        DOWN    2          2    5
1       Enabled 100      2004:1000:1000:1000::10
UP      2        2
```

History

This command was first available in ExtremeXOS 12.1.

Support for IPv6 flow-redirection policies was added in ExtremeXOS 12.7.

Ping success count information added in ExtremeXOS 22.7.

Platform Availability

This command is available for IPv4 and IPv6 flow-redirection policies on the platforms listed for the Policy Based Routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document..

show flowmon

```
show flowmon
```

Description

Displays the state of Flow Monitor.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to show whether Flow Monitor is in enabled or disabled state, the number of groups, number of keys, and number of collectors. It also shows the current total number of flows learned in each hardware domain.

Example

The following command shows Flow Monitor information:

```
# show flowmon
```

History

This command was first available in ExtremeXOS 32.2.

Platform Availability

This command is available on ExtremeSwitching 5520 and 5420 series switches.

show flowmon collector

```
show flowmon collector [collector_name | all] detail
```

Description

Displays Flow Monitor collector information.

Syntax Description

collector	Specifies to send flow information to a collector.
<i>collector_name</i>	Specifies the name of the created collector. Range is 32 characters.
all	Specifies to show Flow Monitor information for all collectors.
detail	Specifies to show all information associated with the specified collector.

Default

N/A.

Usage Guidelines

If **detail** is entered, the source and destination IPv4 addresses, source and destination UDP ports, VR name, and export MTU size are displayed.

If *collector_name* is not entered, then a list of collectors is displayed.

Example

The following command displays detail information for a collector with the name 'c1':

```
# show flowmon collector c1 detail
Name                               : c1
Destination IP address             : 11.1.1.100
Virtual Router                     : VR-Default (default)
Source IP address                  : 11.1.1.1
Source UDP port                    : 46562 (default)
Destination UDP port               : 4739 (default)
Export MTU                         : 512 (default)
Refresh time                       : 600 (default)
* (pacman debug) 5420F-24T-4XE-SwitchEngine.13 #

* (pacman debug) 5420F-24T-4XE-SwitchEngine.13 # show flowmon key k1 detail
Name                               : k1
Type                               : IPv4
Group                              : g1
Destination IP address             : 11.1.1.101
```

History

This command was first available in ExtremeXOS 32.2.

Platform Availability

This command is available on ExtremeSwitching 5520 and 5420 series switches.

show flowmon group

```
show flowmon group group_name | all detail
```

Description

Displays Flow Monitor group information.

Syntax Description

group	Specifies the Flow Monitor group.
<i>group_name</i>	Specifies the assigned name of the Flow Monitor group. Range is 32 characters.
all	Specifies to show Flow Monitor information for all groups.
detail	Specifies to show all information associated with the specified group.

Default

N/A.

Usage Guidelines

For each Flow Monitor group, the group name and parameters are shown.

When **detail** is entered, the assigned collector name, complete template name, template key portion name, and a list of key names, dependent keys and containing groups are shown. In addition, the group state and whether the group is enabled or disabled are shown. The group state is defined as follows:

- Configuring – The user is configuring the group, and the group is not enabled.
- GroupInstall – The group is being installed.
- KMirrorInstall – The K-mirror rule is being installed for the group.
- KeysInstall – The rules related to the keys that are added to the group are being installed.
- Ready – The group is fully installed and awaiting hardware activation.
- CollectorActivate – The collector is being informed of the complete templates associated to the group.
- GroupActivate – The group is being activated.
- Active – The group is fully active, and flows are being collected and advertised to the collector.

When **all** is entered, then all groups are shown with full detail.

Example

The following command displays detail information for the group 'g1':

```
# show flowmon group g1 detail
Name                               : g1
Group is configured to be          : Enabled
Group is operationally              : Enabled
Max Flow Age                       : 60000      (default)
Flow Limit                         : 256
Collector                          : c1
Group State                        : Active
Group Last Failed State             : <none>
Group enable backoff in progress    : No
Complete template name              : g1
Key template name                   : k1
Keys added to this group:
  k1
```

History

This command was first available in ExtremeXOS 32.2.

Platform Availability

This command is available on ExtremeSwitching 5520 and 5420 series switches.

show flowmon group statistics

```
show flowmon group group_name statistics no-refresh wide
```

Description

Displays Flow Monitor group statistics information.

Syntax Description

group	Specifies the Flow Monitor group.
<i>group_name</i>	Specifies the assigned name of the Flow Monitor group. Range is 32 characters.
statistics	Specifies to show Flow Monitor information for all groups.
no-refresh	Specifies to show only one sample of the statistics and then exit the statistics display.
wide	Specifies to show a wider version of the display.

Default

N/A.

Usage Guidelines

When *group_name* is entered, only the statistics for that group are shown. When not entered, all groups are shown using the **Page Up** and **Page Down** options.

Each group name is displayed along with the number of active flows, the number of flows that have aged out, the number of flows that were not learned because the group flow limit was exceeded, the number of flows learned on the group, and the number of flows that were not learned because the flow table was full.

When **no-refresh** is provided, the specified group or groups are displayed one time with their current values.

If **no-refresh** is not entered, the display is updated with refreshed values every five seconds. This display operates as a standard updatable Switch Engine statistics display, providing **Page Up**, **Page Down**, **Exit**, and **Clear** options.

When **no-refresh** is entered, the specified group or groups are displayed one time with their current values.

When **wide** is entered, a wider version of the display is shown. When **wide** is not entered, the counters are limited to a display width that fits into 80 display columns. When **wide** is provided, the full 64-bit counter value can be displayed (about 21 columns per integer value).

Example

The following command shows statistics for a group with the name 'g1' with **no-refresh** specified:

```
# show flowmon gr g1 statistics no-refresh
Group Name           Active   Aged Out Flow Limit   Learned Full Table State
                   Flows      Flows      Drops      Flows      Drops
-----
g1                   1         0         0         1         0 A
-----
> in field indicates value exceeds column width. Use 'wide' option or '0' to clear.
State: A-Active, R-Ready, NR-Not Ready
```

History

This command was first available in ExtremeXOS 32.2.

Platform Availability

This command is available on ExtremeSwitching 5520 and 5420 series switches.

show flowmon group template

```
show flowmon group group_name template detail
```

Description

Displays Flow Monitor group template information.

Syntax Description

group	Specifies the Flow Monitor group.
<i>group_name</i>	Specifies the assigned name of the Flow Monitor group. Range is 32 characters.
template	Specifies to show template information for the specified group.
detail	Specifies to show all information associated with the specified group.

Default

N/A.

Usage Guidelines

The complete template name is shown with the template key portion name.

When **detail** is entered, the Element IDs are shown in text form.

Example

The following command shows template information for the group 'max-flow-age':

```
# show flowmon group max-flow-age template
```

History

This command was first available in ExtremeXOS 32.2.

Platform Availability

This command is available on ExtremeSwitching 5520 and 5420 series switches.

show flowmon key

```
show flowmon key key_name
```

Description

Displays Flow Monitor key information.

Syntax Description

key	Specifies the Flow Monitor key.
<i>key_name</i>	Specifies the assigned name of the Flow Monitor key. Range is 32 characters.
all	Specifies to show Flow Monitor information for all keys.
detail	Specifies to show all information associated with the specified key.

Default

N/A.

Usage Guidelines

For each key, the name, type, associated group name (if applicable), and the 'before' or after key name (if applicable) are displayed.

When **all** is entered and **detail** is not specified, all keys are shown in tabular format.

When **detail** is entered, the IPv4 or IPv6 address and mask, source or destination UDP or TCP port, protocol or next header, and specified list of ports to which the key is applied are shown. If none are specified, all values are shown.

Example

The following command displays information for a key with the name 'src-ipv4-addr':

```
# show flowmon key src-ipv4-addr
```

History

This command was first available in ExtremeXOS 32.2.

Platform Availability

This command is available on ExtremeSwitching 5520 and 5420 series switches.

show forwarding configuration

```
show forwarding configuration
```

Description

Displays the configured selection criteria for *ECMP* routes and load-sharing group ports and the hardware table settings, including the configured and current hash algorithm and dual-hash settings.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

The output of this command displays the following information:

- Configured hash algorithm—The hash algorithm configured on the switch. After the configuration is saved and the switch is rebooted, the switch uses this hash algorithm.
- Current hash algorithm—The hash algorithm currently used by the switch.
- Configured dual-hash setting—Whether the dual-hash feature is configured 'on' or 'off' on the switch. After the configuration is saved and the switch is rebooted, the switch uses this setting.
- Current dual-hash setting—Whether the dual-hash feature is currently 'on' or 'off' on the switch.
- Dual-Hash recursion level—The current dual-hash recursion level; default is '1.'
- Sharing criteria—Current selection criterion used for ECMP route sharing as well as for load-sharing groups. Specifies which Layer 3 and Layer 4 information is used in the sharing hash algorithm. For more information, see the description for the [configure forwarding sharing \[L3 | L3_L4\]](#) command.
- Group Table Compression—Whether the group table compression is currently 'on' or 'off' on the switch.
- Switching mode—Whether the switching mode is currently set to 'cut-through' or 'store-and-forward.'
- Fabric flow control—Whether flow control fabric configuration is set to 'auto' or 'off.'
- ARP and ND Settings:—Suppression filter status.
- Extended Edge Switching—VLAN port membership and IP multicast replication status.
- Forward to VLAN filters for mDNS, LLMNR, and UPnP.

It is possible for the values of the configured and the current hash, or the configured and current dual-hash settings to be different. For example, if you modified the hash algorithm and have not saved the configuration and rebooted the switch, the values might be different. In this situation, the switch also displays the following message:

NOTE: A save and reboot are required before the configured hash will take effect

Example

The following command displays the hardware forwarding algorithm configured on the switch:

```
show forwarding configuration

L2 and L3 Forwarding table hash algorithm:
  Configured hash algorithm:      crc32
  Current hash algorithm:         crc32

L3 Dual-Hash configuration:
  Configured setting:             on
  Current setting:                on
  Dual-Hash Recursion Level:      1

Hash criteria for IP unicast traffic for L2 load sharing and ECMP route sharing
  Sharing criteria:               L3_L4

IP multicast:
  Group Table Compression:        on
  Local Network Forwarding:       slow-path
  Lookup-Key:                     (SourceIP, GroupIP, VlanId)
```

```

Internal lookup tables:
  Configured Setting:      12-and-13
  Current Setting:        12-and-13

Switch Settings:
  Switching mode:          store-and-forward

L2 Protocol:
  Fast convergence:        on

Rate Limit:
  Overhead Bytes:          20

Fabric Flow Control:
  Fabric Flow Control:     auto

ARP and ND Settings:
  ARP Suppression Filters: per-port
  ND Suppression Filters:  per-port

VPEX Settings:
  VLAN Port Membership:    hash-table
  IPMC Replication:        bpe

IP multicast:

mDNS: flood

LLMNR: learn

UPnP: flood

```

History

This command was first available in ExtremeXOS 11.3.2.

The flow control feature was added in ExtremeXOS 12.5.

The Forwarding Lookup-Key feaure show output is added in ExtremeXOS 15.3.

Extended Edge Switching setting for VLAN port membership was added in ExtremeXOS 22.6.

Extended Edge Switching IP multicast replication feature information was added in ExtremeXOS 30.6.

Forward to VLAN filters for mDNS, LLMNR, and UPnP was added in Release 32.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show forwarding hardware-utilization

```
show forwarding hardware-utilization {slot [slot_num | all]}
```

Description

Shows forwarding hardware table utilization statistics.

Syntax Description

hardware-utilization	Specifies showing hardware table utilization statistics.
slot	Specifies selecting slots.
<i>slot_num</i>	Specifies which slots to show statistics for.
all	Specifies showing statistics for all slots.

Usage Guidelines

Shows forwarding hardware table utilization statistics with L2, L3 and ACLs configured.

Example

The following example shows table utilization statistics:

```
# show forwarding hardware-utilization

Hardware Table Utilization Statistics

Slot: 1
Type: X870-32c
Resource Type          Current    Maximum    % Util.
-----
MAC Entries            0         73727      0
Host Entries           0         73728      0
IPv4 Entries           0         73728      0
IPv6 Entries           0        36864      0
Long IPv6 Entries      0         2048       0
Total Routes           0        16384      0
IPv4 Neighbors         0          N/A       N/A
IPv6 Neighbors         0          N/A       N/A
IPv4 Routes            0          N/A       N/A
IPv6 Routes            0          N/A       N/A
ECMP Next Hops         0         1023       0
ACL Ingress Entries    40         6144       0
ACL Ingress Counters   0        38912      0
ACL Ingress Meters     0         2048       0
ACL Ingress Slices     0          12       0
ACL Egress Entries     0         1024       0
ACL Egress Counters    0         1024       0
ACL Egress Meters      0         1024       0
ACL Egress Slices      0          4       0
ACL VLAN Entries *      1         2048       0
ACL VLAN Slices *      1          4       25

Legend: N/A - Maximum not defined in published sFlow structure.
* - Resource type not defined in published sFlow structure.
    Also see IP usage in 'show iproute reserved-entries statistics'.
```

History

This command was first available in ExtremeXOS 22.6.

Legend explaining N/A entries was added in ExtremeXOS 30.1.

ACL VLAN entries and slices was added in ExtremeXOS 30.7.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show heartbeat process

```
show heartbeat process {name}
```

Description

Displays the health of the ExtremeXOS processes.

Syntax Description

<i>name</i>	Specifies the name of the process.
-------------	------------------------------------

Default

N/A.

Usage Guidelines

The software monitors all of the ExtremeXOS processes running on the switch. This process monitor creates and terminates XOS processes on demand (for example, when you log in or log out of the switch) and restarts processes if an abnormal termination occurs (for example, if your system crashes). The process monitor also ensures that only version-compatible processes and processes with proper licenses are started.

The `show heartbeat process` command is a resource for providing background system health information because you can view the health of ExtremeXOS processes on the switch.

Use this command to monitor the health of the ExtremeXOS processes. The switch uses two algorithms to collect process health information: polling and reporting. Both polling and reporting measure the heartbeat of the process. Polling occurs when a HELLO message is sent and a HELLO_ACK message is received. The two counts are the same. Reporting occurs when a HELLO_ACK message is sent only. Therefore, no HELLO messages are sent and the HELLO count remains at zero.

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

- Process Name—The name of the process.
- Hello—The number of hello messages sent to the process.
- HelloAck—The number of hello acknowledgement messages received by the process manager.
- Last Heartbeat Time—The timestamp of the last health check received by the process manager. (Unknown specifies kernel modules and they do not participate in heartbeat monitoring.)

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

You may find it useful to capture the process information under normal operating conditions to establish a baseline. By having a baseline, if you experience a problem, you and your technical support representative can more easily identify the problem.

Example

To display the health of all processes on your system, use the following command:

```
show heartbeat process
```

The following is sample output from a switch:

Process Name	Hello	HelloAck	Last Heartbeat Time
aaa	0	254328	Tue Feb 10 05:21:46 2004
acl	50867	50867	Tue Feb 10 05:21:43 2004
bgp	0	0	Wed Feb 4 08:03:18 2004
cfgmgr	25433	25433	Tue Feb 10 05:21:33 2004
cli	84779	84779	Tue Feb 10 05:21:47 2004
cna	20234	20234	Mon Feb 9 00:28:35 2004
devmgr	0	250507	Tue Feb 10 05:21:47 2004
dirser	0	254336	Wed Feb 4 08:03:18 2004
dosprotect	0	254335	Tue Feb 10 05:21:47 2004
eaps	0	254336	Tue Feb 10 05:21:48 2004
edp	50867	50867	Tue Feb 10 05:21:44 2004
elrp	50867	50867	Tue Feb 10 05:21:43 2004
ems	63584	63584	Tue Feb 10 05:21:44 2004
epm	0	0	Wed Feb 4 08:03:18 2004
esrp	50867	50867	Tue Feb 10 05:21:46 2004
...			

To display the health of the *STP* process on your system, use the following command:

```
show heartbeat process stp
```

The following is sample output from a switch:

Process Name	Hello	HelloAck	Last Heartbeat Time
stp	50870	50870	Tue Feb 10 05:22:13 2004

History

This command was first available in an ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show identity-management blacklist

```
show identity-management blacklist
```

Description

Displays the identities in the identity manager blacklist.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays the identities in the blacklist:

```
* Switch.93 # show identity-management blacklist
-----
Type      BlackList Entry
-----
MAC       01:02:03:04:05:06/ff:ff:ff:00:00:00
IP        1.2.3.4/255.255.255.0
User      john@mydomain.com
-----
> indicates entry value truncated past 35 characters
Number of BlackList Entries      : 3
```

History

This command was first available in ExtremeXOS 12.7.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show identity-management entries

```
show identity-management entries {user id_name} {domain domain} {ports
    port_list} {mac mac_address} {vlan vlan_name} {ipaddress ip_address}
    {detail}
```

Description

Displays the entries in the identity management database.

Syntax Description

<i>id_name</i>	Limits the display to entries that contain the specified user ID.
<i>domain</i>	Limits the display to entries for the specified domain.
<i>port_list</i>	Limits the display to entries for the specified ports.
<i>mac_address</i>	Limits the display to entries that contain the specified MAC address.
<i>vlan_name</i>	Limits the display to entries that contain the specified <i>VLAN</i> name.
<i>ip_address</i>	Limits the display to entries that contain the specified IP address.
detail	Expands the display to include more information about identity management entries.

Default

N/A.

Usage Guidelines

Only admin-level users can execute this command.

The displayed ID Name is the actual user name when Network Login or Kerberos Snooping is enabled. For unknown users, the software creates a user name using the format: User_XXXXXXXXXXXXXXXX. The number in the user name is a 16-bit hash number that is generated using the user's port, MAC address, and IP address numbers.

The displayed Domain Name is displayed only if the client is discovered through Kerberos snooping or Dot1x and the domain name is supplied in the form of *domain\user*). The NetBIOS hostname is only displayed if this information was present in the Kerberos packets.

When the role is shown as multiple, the identity is connected through multiple ports/locations and different roles apply to each device.

Example

The following command displays all entries in the identity management database:

```
* Switch.4 # show identity-management entries
ID Name/      Flags  Port      MAC/      VLAN      Role
Domain Name   IP
-----
Unknown_00:00:00:> ----  1:3      00:00:00:00:00:22  v1(1)      unauthentic
-- NA --
00005A4B0000  -m--  1:4      00:00:5a:4b:d1:98  test126(1)  Phone
126.0.0.2(1)
00005A4B0000  -m--  1:4      00:00:5a:4b:d1:9c  test128(1)  Phone
128.0.0.2(1)
00005A4B0000  -m--  1:4      00:00:5a:4b:d1:9e  test129(1)  Phone
129.0.0.2(1)
.
.
.
000105000000  -m--  1:4      00:01:05:00:03:18  test150(1)  Phone
```

```
-- NA --
OTHER(00:04:96:1e> 1--- 4:11 00:04:96:1e:32:80 -- NA -- unauthentic>
-- NA --
joe --k- 1 00:00:22:33:55:66 v1(1) authenticated
extreme 2.1.3.4(1)
bill --k- 2 00:00:22:33:44:55 v1(2) multiple
corp.extremenetworks.com 1.2.3.4(1)
Unknown_00:00:00:> ---- 1 00:00:00:00:22:33 v1(1) unauthentic>
-- NA --
.
.
.
OTHER(02:04:96:51> 1--- 4:3 02:04:96:51:77:c7 -- NA -- unauthentic>
-- NA --
-----
Flags: k - Kerberos Snooping, l - LLDP Device,
m - NetLogin MAC-Based, w - NetLogin Web-Based,
x - NetLogin 802.1X
Legend: > - VLAN / ID Name / Domain / Role Name truncated to column width
(#) - Total # of associated VLANs/IPs
-- NA --- No IP or VLAN associated
Total number of entries: 60
```

The following command shows the detail format:

```
* Switch.4 # show identity-management entries detail
- ID: "00005A4B0000", 1 Port binding(s)
Role: "Phone"
Port: 1:4, 24 MAC binding(s)
MAC: 00:00:5a:4b:d1:98, Flags: -m--, Discovered: Fri Sep 24 18:30:17 2010
1 VLAN binding(s)
VLAN: "test126", 1 IP binding(s)
IPv4: 126.0.0.2
Security Profile: ----, Security Violations: ----;
MAC: 00:00:5a:4b:d1:9c, Flags: -m--, Discovered: Fri Sep 24 18:30:17 2010
1 VLAN binding(s)
VLAN: "test128", 1 IP binding(s)
IPv4: 128.0.0.2
Security Profile: ----, Security Violations: ----;
MAC: 00:00:5a:4b:d1:9e, Flags: -m--, Discovered: Fri Sep 24 18:30:17 2010
1 VLAN binding(s)
VLAN: "test129", 1 IP binding(s)
IPv4: 129.0.0.2
Security Profile: ----, Security Violations: ----;
.
.
.
MAC: 00:00:5a:4b:d1:c8, Flags: -m--, Discovered: Fri Sep 24 18:30:17 2010
1 VLAN binding(s)
VLAN: "test150", 1 IP binding(s)
IPv4: 150.0.0.2
Security Profile: ----, Security Violations: ----;
- ID: "000071710000", 1 Port binding(s)
Role: "Phone"
Port: 1:5, 1 MAC binding(s)
MAC: 00:00:71:71:00:01, Flags: -m--, Discovered: Fri Sep 24 19:42:29 2010
1 VLAN binding(s)
VLAN: "palani", 0 IP binding(s)
- ID: "000105000000", 1 Port binding(s)
Role: "Phone"
Port: 1:4, 25 MAC binding(s)
MAC: 00:01:05:00:03:00, Flags: -m--, Discovered: Fri Sep 24 18:30:17 2010
1 VLAN binding(s)
```

```

VLAN: "test126", 0 IP binding(s)
MAC: 00:01:05:00:03:01, Flags: -m--, Discovered: Fri Sep 24 18:30:17 2010
1 VLAN binding(s)
VLAN: "test127", 0 IP binding(s)
MAC: 00:01:05:00:03:02, Flags: -m--, Discovered: Fri Sep 24 18:30:17 2010
1 VLAN binding(s)
VLAN: "test128", 0 IP binding(s)
.
.
.
MAC: 00:01:05:00:03:18, Flags: -m--, Discovered: Fri Sep 24 18:30:18 2010
1 VLAN binding(s)
VLAN: "test150", 0 IP binding(s)
- ID: "OTHER(00:04:96:1e:32:80)", 8 Port binding(s)
Role: "unauthenticated"
Port: 4:11, 1 MAC binding(s)
MAC: 00:04:96:1e:32:80, Flags: l---, Discovered: Fri Sep 24 18:30:17 2010
0 VLAN binding(s)
Port: 4:12, 1 MAC binding(s)
MAC: 00:04:96:1e:32:80, Flags: l---, Discovered: Fri Sep 24 18:30:17 2010
0 VLAN binding(s)
Port: 4:13, 1 MAC binding(s)
MAC: 00:04:96:1e:32:80, Flags: l---, Discovered: Fri Sep 24 18:30:17 2010
0 VLAN binding(s)
.
.
.
Port: 4:18, 1 MAC binding(s)
MAC: 00:04:96:1e:32:80, Flags: l---, Discovered: Fri Sep 24 18:30:17 2010
0 VLAN binding(s)
- ID: "OTHER(02:04:96:51:77:c7)", 2 Port binding(s)
Role: "unauthenticated"
Port: 1:1, 1 MAC binding(s)
MAC: 02:04:96:51:77:c7, Flags: l---, Discovered: Fri Sep 24 18:30:17 2010
0 VLAN binding(s)
Port: 4:3, 1 MAC binding(s)
MAC: 02:04:96:51:77:c7, Flags: l---, Discovered: Fri Sep 24 18:30:17 2010
0 VLAN binding(s)
-----
Flags:          k - Kerberos Snooping, l - LLDP Device,
m - NetLogin MAC-Based, w - NetLogin Web-Based,
x - NetLogin 802.1X
Security Profile:  a - ARP Validation, d - DoS Protection,
g - Gratuitous ARP Protection, r - DHCP Snooping
Security Violations: A - ARP Validation Violation, D - DoS Violation
G - Gratuitous ARP Violation, R - Rogue DHCP Server Detected

```

The following command example shows how domain names, NetBIOS hostnames, and multiple roles appear when in use:

```

Switch.4 # show identity-management entries detail
- ID: "john", 1 Port binding(s)
Role: "IT-Engineer"
Domain: "XYZCorp.com", NetBios hostname: "JOHN-DESKTOP"
Port: 17 (Bld-1-Port-1), 1 MAC binding(s)
MAC: 00:00:5a:4b:d1:98, Flags: --k-, Discovered: Tue Nov 16 12:22:46 2010
Force Aging TTL: 00:00:02    Inactive Aging TTL: 00:00:03
1 VLAN binding(s)
VLAN: "corp", 1 IP binding(s)
IPv4: 126.0.0.2
Security Profile: -d--, Security Violations: ----;
- ID: "ramesh", 2 Port binding(s)
Role: "multiple"

```

```

Domain: "corp.extremenetworks.com"
Port: 1, 1 MAC binding(s)
MAC: 00:00:00:00:00:13, Flags: --k-, Discovered: Sat Apr  2 02:23:41 2011
Force Aging TTL: 00:00:02      Inactive Aging TTL: N/A
1 VLAN binding(s)
VLAN: "v1", 1 IP binding(s)
IPv4: 10.120.89.9
Role: "Engineer"
Security Profile: adgsr---, Security Violations: A-----,
Port: 2, 1 MAC binding(s)
MAC: 00:00:00:00:00:30, Flags: --k-, Discovered: Sat Apr  2 02:24:30 2011
Force Aging TTL: 00:00:02      Inactive Aging TTL: N/A
1 VLAN binding(s)
VLAN: "v2", 1 IP binding(s)
IPv4: 10.2.3.45
Role: "iphoneEngineer"
Security Profile: ----, Security Violations: ----;
-----
Flags:                k - Kerberos Snooping, l - LLDP Device,
m - NetLogin MAC-Based, w - NetLogin Web-Based,
x - NetLogin 802.1X
Security Profile:      a - ARP Validation, d - DoS Protection,
g - Gratuitous ARP Protection, r - DHCP Snooping
Security Violations: A - ARP Validation Violation, D - DoS Violation
G - Gratuitous ARP Violation, R - Rogue DHCP Server Detected

```

The following command example shows that you can specify multiple options, such as the user name and ports:

```
show identity-management entries user eelco ports 2:2
```

History

This command was first available in ExtremeXOS 12.4.

Kerberos Force Aging TTL and Inactive Aging TTL information was added in ExtremeXOS 12.6.

Support for multiple roles for a single identity was added in ExtremeXOS 12.7.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show identity-management greylist

```
show identity-management greylist
```

Description

Displays the identities in the identity manager greylist.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to display the identities in the identity manager greylist.

Example

```
* Switch.94 # show identity-management greylist
-----
Type      GreyList Entry
-----
User      june@mydomain.com
User      Richard@corp.acme.com
-----
> indicates entry value truncated past 35 characters
Number of GreyList Entries      : 2
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show identity-management list-precedence

```
show identity-management list-precedence
```

Description

This command displays the order of list-precedence. The default list-precedence is: greylist blacklist whitelist.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to display the order of list-precedence.

Example

```
* Switch.97 # show identity-management list-precedence
List Precedence:
1.    Greylist
2.    Blacklist
3.    Whitelist
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show identity-management role

```
show identity-management role {role_name} {detail}
```

Description

Displays summary or detailed configuration information for one or all roles.

Syntax Description

<i>role_name</i>	Specifies a name of an existing role to display.
all	Specifies that all roles are to be displayed.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays all roles that are configured on the switch:

```
* Switch.95 # show identity-management role
-----
Role Name          Priority    Child Roles    # Identities
*authenticated     255                0
*unauthenticated   255                0
extr-empl          255      extr-engr      2
extr-engr          255                0
*whitelist         0                0
*blacklist         0                3
```

```
-----
Flags : * - Default Roles
-----
Total number of role(s) configured : 6
```

The following command displays detailed information for all roles that are configured on the switch:

```
* Switch.96 # show identity-management role detail
Role name : extr-empl
Child Roles : engr
Match Criteria : "company==Extreme;"
Policies : extrPol
Identities : john_smith@d.com; MAC: 00:16:23:51:77:99; Port:8
bob_craig@e.com; MAC: 00:18:23:51:77:99; Port:9
Role name : engr
Child Roles : india-engr
Match Criteria : "department==Engineering;"
Policies : engrPol, extrPol
Identities : joe_hardy@b.com; MAC: 00:12:23:51:77:99; Port:10
Role name : india-engr
Child Roles : -
Match Criteria : "country=India; AND department=Engineering;"
Policies : indEngrPol, engrPol, extrPol
Identities : bill_jacob@b.com; MAC: 00:12:33:51:77:99; Port:11
Role name : marketing
Child Roles : -
Match Criteria : "department=Marketing;"
Policies : markrPol, extrPol
Identities : will_smith@a.com; MAC: 00:11:33:51:77:99; Port:14
Role Name: whitelist (Default Role)
Child Roles : ---
Priority : 0
Match Criteria : "Not Applicable"
Policies : --
Identities # : 0
Identities : --
Role Name: blacklist(Default Role)
Child Roles : ---
Priority : 0
Match Criteria : "Not Applicable"
Policies : --
Identities # : 3
Identities : Unknown_00:11:22:33:44:55; MAC: 00:11:22:33:44:55; Port:1
johndoe@extremenetworks.com; MAC: 00:01:02:03:04:05; Port:2
janedoe@extremenetworks.com; MAC: 00:02:04:06:08:10; Port:3
```

The next two examples display detailed information for a single role:

```
* Switch.97 # show identity-management role extr-empl detail
Role name : extr-empl
Child Roles : engr
Match Criteria : "company=Extreme;"
Policies : extrPol
Identities : johnsmith@extreme.com; MAC: 00:11:33:55:77:99; Port:4
bobcraig@extreme.com; MAC: 00:01:03:05:07:09; Port:5
* Switch.98 # show identity-management role NotAccessibleUser detail
Role name : NotAccessibleUser
Child Roles : engr
Match Criteria : "UserName = adam; AND IP-Address == 1.2.3.0/24; AND port == 1;"
Policies : extrPol
Identities : adam; MAC: 00:00:11:22:33:44; Port: 1
```

History

This command was first available in ExtremeXOS 12.5.

MAC addresses were added to the displays for the detail option in ExtremeXOS 12.7.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show identity-management statistics

```
show identity-management statistics
```

Description

Displays operation statistics for the identity management feature.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

A user can login from multiple machines or ports, so the total number of login instances can be more than the number of unique users logged in.

Example

The following command displays the identity management feature statistics:

```
Switch.4 # show identity-management statistics
Total number of users logged in      :    2
Total number of login instances     :    2
Total memory used                    :    1 Kbytes
Total memory used by events          :    0 Kbytes
Total memory available               :   511 Kbytes
High memory usage level reached count :    0
Critical memory usage level reached count:    0
Max memory usage level reached count  :    0
Current memory usage level           : Normal
Normal memory usage level trap sent   :    0
High memory usage level trap sent     :    0
Critical memory usage level trap sent :    0
Max memory usage level trap sent      :    0
Event notification sent               :    0
Total number of roles configured      :    3
```


History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show identity-management whitelist

```
show identity-management whitelist
```

Description

Displays the identities in the identity manager whitelist.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays the identities in the whitelist:

```
* Switch.94 # show identity-management whitelist
-----
Type      WhiteList Entry
-----
MAC       04:32:13:44:25:06/ff:ff:ff:00:00:00
IP        11.12.13.14/255.255.255.0
User      jane@mydomain.com
-----
> indicates entry value truncated past 35 characters
Number of WhiteList Entries      : 3
```

History

This command was first available in ExtremeXOS 12.7.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show identity-management

show identity-management

Description

Displays the identity management feature configuration.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays the identity management feature configuration:

```
# show identity-management
Identity Management           : Enabled
Stale entry age out (effective) : 180 Seconds (180 Seconds)
Max memory size               : 512 Kbytes
Enabled ports                 : 1-26
FDB Detection Disabled ports  : None
IPARP Detection Disabled ports : None
IPSecurity Detection Disabled ports : None
Kerberos Detection Disabled ports : None
LLDP Detection Disabled ports  : None
Netlogin Detection Disabled ports : None
SNMP trap notification        : Enabled
Match Criteria Inheritance     : On
Access list source address type : MAC
Kerberos aging time (DD:HH:MM) : None
Kerberos force aging time (DD:HH:MM) : None
Valid Kerberos servers         : none configured(all valid)
```

History

This command was first available in ExtremeXOS 12.4.

Kerberos Force Aging Time information was added in ExtremeXOS 12.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show igmp

```
show igmp {vlan} {vlanname}
```

Description

This command can be used to display an *IGMP*-related configuration and group information, per *VLAN*.

Syntax Description

<i>vlanname</i>	Specifies a VLAN name.
-----------------	------------------------

Default

N/A.

Usage Guidelines

The output of this command shows:

- The VLAN name.
- The router interface IP address and subnet mask.
- If the interface is active (up), by the letter U.
- If IP forwarding is enabled for the interface, by the letter f.
- If multicast forwarding is enabled, by the letter M.
- If IGMP is enabled, by the letter i.
- If IGMP snooping is enabled, by the letter z.

Example

The following command displays the IGMP configuration:

```
# show igmp
VLAN          IP Address      Flags      nLRMA  nLeMA  IGMPver
Default       0.0.0.0         / 0  ---izpt-    0      0      3
isc           50.50.50.1      /24  ---izpt-    0      0      3
v1            0.0.0.0         / 0  U---izpt-    0      2      3
v3000         1.1.1.1         /24  ---izpt-    0      0      3
v666          6.0.0.1         /16  ---izpt-    0      0      3
Flags: (f) Forwarding Enabled, (i) IGMP Enabled
(m) Multicast Forwarding Enabled, (p) IGMP Proxy Query Enabled
(r) Receive Router Alert Required (t) Transmit Router Alert
(U) Interface Up, (z) IGMP Snooping Enabled
(nLeMA) Number of Learned Multicast Addresses
(nLRMA) Number of Locally Registered Multicast Addresses
```

The following command displays the IGMP configuration for VLAN v1:

```
# show igmp v1
Query Interval      : 125 sec
Max Response Time   : 10 sec
Last Member Query   : 1 sec
Robustness          : 2
```

```

Interface on VLAN vlan1 is enabled and up.
inet 0.0.0.0/0
Locally registered multicast addresses:
Learned multicast addresses (Last Querier=118.1.1.100):
224.0.0.2          224.0.0.22
s = static igmp member
Flags:
IP Fwding  NO          IPmc Fwding  NO          IGMP YES
IGMP Ver   V3          Snooping   YES          Proxy Query YES
XmitRtrAlrt YES       RcvRtrAlrtReq NO

```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IPv4 multicast feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

show igmp counters

```
show igmp counters
```

Description

Shows Internet Group Management Protocol (IGMP) counters.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Example

The following example shows IGMP counters:

```

# show igmp counters
Global IGMP Statistics:
Out Query   =          13 Out Report =          0 Out Leave =          0
In Query    =          0 In Report  =          13 In Leave  =          0
In Error    =          0

```

History

This command was first available in ExtremeXOS 30.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IPv4 multicast feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

show igmp group

```
show igmp group {{vlan} {name} | {grpipaddress}} {IGMPv3}
```

Description

Lists the *IGMP* group membership for the specified *VLAN*.

Syntax Description

<i>name</i>	Specifies a VLAN name.
<i>grpipaddress</i>	Specifies a group IP address.
IGMPv3	Displays the IGMP group in IGMPv3 format (if group record is IGMPv3 compatible, otherwise displays in earlier format).

Default

IGMPv2.

Usage Guidelines

If no VLAN is specified, all VLANs are displayed. You can also filter the display by group address or by multicast stream sender address.

The output of this command shows:

- The multicast group address received.
- The version of the IGMP group.
- The name of the VLAN where the group address is being received.
- The physical port where the group address is being received. If multiple ports within the VLAN have subscribers for the group, all the ports are listed.
- The age since the last IGMP report for this group was received.



Note

The `show igmp group` command output is populated on the router that is the PIM Rendezvous Point.

Example

The following is sample output from the `show igmp group` command, listing the IGMP group membership:

Group Address	Ver	Vlan	Port	Age
239.2.4.70	2	banana	7	101
224.0.1.24	2	banana	7	107
239.255.255.254	2	banana	7	103
Total: 3				

History

This command was first available in ExtremeXOS 10.1.

The **IGMPv3** option was added in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IPv4 multicast feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

show igmp snooping cache

This command is provided for backward compatibility. The recommended command is:

```
show mcast cache {{vlan} name} {[group grpaddressMask | grpaddressMask]
  {source sourceIP | sourceIP}} {type [snooping | pim | mvr]}|
  {summary}}
```

The syntax for the original form of this command is:

```
show igmp snooping cache {{vlan} name} {{group} grpaddressMask}
```

Description

Displays multicast cache entries added by *IGMP* snooping for all VLANs and groups. The display can be limited to specific VLANs or groups.

Syntax Description

<i>name</i>	Specifies a <i>VLAN</i> name.
<i>grpaddressMask</i>	Specifies a multicast group address and mask.

Default

Displays information for all VLANs and groups.

Usage Guidelines

None.

Example

The following command displays IGMP snooping cache information for all VLANs and groups:

```
show igmp snooping cache
```

This command display is the same as for the following preferred command:

```
show mcast cache {{vlan} name} {[group grpaddressMask | grpaddressMask]
{source sourceIP | sourceIP}} {type [snooping | pim | mvr]}| {summary}}
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IGMP snooping feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

show igmp snooping vlan filter

```
show igmp snooping {vlan} name filter
```

Description

Displays *IGMP* snooping filters.

Syntax Description

<i>name</i>	Specifies a <i>VLAN</i> name.
-------------	-------------------------------

Default

None.

Usage Guidelines

Use this command to display IGMP snooping filters configured on the specified VLAN. When no VLAN is specified, all the filters are displayed.

Example

The following command displays the IGMP snooping filter configured on VLAN vlan101:

```
# show igmp snooping vlan101 filter
Filter          Port Flags
igmppermit0    5:10 a
Flags: (a) Active
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IGMP snooping feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

show igmp snooping vlan static

```
show igmp snooping {vlan} name static [group | router]
```

Description

Displays static *IGMP* snooping entries.

Syntax Description

<i>name</i>	Specifies a <i>VLAN</i> name.
group	Displays static multicast groups.
router	Displays static router entries.

Default

None.

Usage Guidelines

Use this command to display the IGMP snooping static groups or router ports configured on the specified VLAN. When no VLAN is specified, all the static groups or router ports are displayed.

Example

The following command displays the IGMP snooping static groups configured on VLAN vlan101:

```
# show igmp snooping vlan101 static group
VLAN vlan101 (4094)
Group      Port  Flags
239.1.1.2  29    s-
```



```
239.1.1.2      30      s-
239.1.1.2      31      sa
239.1.1.2      32      s-
239.1.1.2      34      s-
Total number of configured static IGMP groups = 5
Flags: (s) Static, (a) Active
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IGMP snooping feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

show igmp snooping vlan

```
show igmp snooping {vlan} name {port port} {IGMPv3}
```

Description

Displays *IGMP* snooping registration information for a specific *VLAN*. The display can be further limited to a specific port or to only IGMPv3 entries.

Syntax Description

<i>name</i>	Specifies a VLAN name.
<i>port</i>	Specifies a single port for which information is displayed.
IGMPv3	Displays the IGMP group in IGMPv3 format (if group record is IGMPv3 compatible, otherwise displays in earlier format).

Default

IGMPv2.

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 membership information
- Router entry
- Timeout information
- Sender entry

Example

The following output displays IGMP snooping registration information on VLAN v1:

```
# show igmp snooping vlan v1
Router Timeout      :      260 sec
Host Timeout       :      260 sec
Igmp Snooping Fast Leave Time : 1000 ms
VLAN v1 d          (4084) Snooping=Enabled
Port   Host        Subscribed   Age   Group-Limit
25     118.1.1.100   All Groups    3     0
```

The following command displays IGMP snooping registration information for port 2:1 on VLAN test:

```
show igmp snooping test port 2:1
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IGMP snooping feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

show igmp snooping

```
show igmp snooping {detail {IGMPv3}}
```

Description

Displays *IGMP* snooping registration information for all VLANs.

Syntax Description

detail	Displays the information in detailed format.
IGMPv3	Displays the IGMP group in IGMPv3 format (if group record is IGMPv3 compatible, otherwise displays in earlier format).

Default

IGMPv2.

Usage Guidelines

None.

Example

The following command displays IGMP snooping registration information for all VLANs:

```
# show igmp snooping
Igmp Snooping Flag      : forward-all-router
Igmp Snooping Flood-list : none
Igmp Snooping Proxy     : Disable
Igmp Snooping Filters   : per-port
Vlan      Vid  Port  #Senders #Receivers Router Enable
-----
Default   1    0      0           Yes
v1        4090  0      0           Yes
```

History

This command was first available in ExtremeXOS 10.1.

The IGMP Forwarding Lookup mode output was removed from this command in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IGMP snooping feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

show igmp ssm-map

```
show igmp ssm-map {group_ip} {vr vr-name}
```

Description

Displays the *IGMP* SSM feature status (enabled or disabled), the mappings for the specified multicast group IP address, and the total count of maps.

Syntax Description

<i>group_ip</i>	Specifies an IP multicast group, for which all mappings in the PIM SSM range are to be displayed. If no group address is specified, the switch displays all IGMP-SSM mappings.
<i>vr-name</i>	Specifies a virtual router name. If the VR name is omitted, the switch displays the mappings on the VR specified by the current CLI VR context.

Default

N/A.

Usage Guidelines

When a target group is specified, this command displays all mapping entries for the configured range in which the group IP address resides.

Example

The following example displays the mappings for the multicast group starting at IP address 232.1.1.1:

```
X620-16x-DUT1.1 # show igmp ssm-map
IGMP SSM mapping : Enabled

Group                Source          Type
232.1.1.1/32         2.1.1.3         Static
232.1.1.2/32         2.1.1.3         Static
232.1.1.3/32         2.1.1.3         Static
232.1.1.4/32         2.1.1.3         Static
232.1.1.5/32         2.1.1.3         Static
232.1.1.6/32         2.1.1.3         Static
232.1.1.7/32         2.1.1.3         Static
232.1.1.8/32         2.1.1.3         Static
232.1.1.9/32         2.1.1.3         Static
232.1.1.10/32        2.1.1.3         Static
232.1.1.11/32        2.1.1.3         Static
232.1.1.12/32        2.1.1.3         Static
232.1.1.13/32        2.1.1.3         Static
232.1.1.14/32        2.1.1.3         Static
232.1.1.15/32        2.1.1.3         Static
232.1.1.16/32        2.1.1.3         Static
232.1.1.17/32        2.1.1.3         Static
232.1.1.18/32        2.1.1.3         Static
232.1.1.19/32        2.1.1.3         Static
232.1.1.20/32        2.1.1.3         Static
232.1.1.21/32        2.1.1.3         Static
Total Number of Mappings : 510
X620-16x-DUT1.2 # show igm ssm-map 232.1.1.1
IGMP SSM mapping : Enabled
Group   : 232.1.1.1
Type    : STATIC
DNS Name: -
Aging TTL (Expiry Time): -
Sources :
        2.1.1.3
Total Number of Mappings : 1
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IPv4 multicast feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

show inline-power configuration ports

```
show inline-power configuration ports {port_list}
```

Description

Displays inline power configuration information for the specified ports.

Syntax Description

<i>port_list</i>	Specifies one or more ports. If you do not specify ports, information for all ports appears.
------------------	--

Default

N/A.

Usage Guidelines

The output displays the following inline power configuration information for the specified ports:

- Config—Indicates whether the port is enabled to provide inline power:
 - Enabled: The port can provide inline power.
 - Disabled: The port cannot provide inline power.
- Operator Limit—Displays the configured limit, in milliwatts, for inline power on the port.
- Label—Displays a text string, if any, associated with the port.

Example

The following command displays inline power configuration information for ports 1 to 7:

```
# show inline-power configuration ports 1-7
```

Port	Config	Operator	Priority	Detection	Classification
Fast	Label	Limit			
1	Enabled	15400 mW	Low	802.3af-only	802.3af-high
N					
2	Enabled	90000 mW	Low	802.3af-only	802.3af
N					
3	Enabled	90000 mW	Low	802.3af-only	802.3af
N					
4	Enabled	90000 mW	Low	802.3af-only	802.3af
N					
5	Enabled	90000 mW	Low	802.3af-only	802.3pre-at
N					
6	Enabled	90000 mW	Low	802.3af-only	802.3pre-bt
N					
7	Enabled	90000 mW	Low	802.3af-only	802.3bt-type4
N					

The following command displays inline power information with **class-based** configured for ports 1 and 2:

```
# show inline-power configuration ports 1 2 3
```

Port Label	Config	Operator Limit	Priority	Detection	Classification	Fast
1	Enabled	class-based	Low	802.3af-only	802.3bt-type4	
2	Enabled	class-based	Low	802.3af-only	802.3bt-type4	
3	Enabled	90000 mW	Low	802.3af-only	802.3bt-type4	

History

This command was first available in ExtremeXOS 11.1.

Per-port Fast PoE information was added in ExtremeXOS 31.1.

The **802.3af**, **802.3af-high**, and **802.3pre-at** options were added in ExtremeXOS 31.7.

class-based information was added in ExtremeXOS 32.3.

Platform Availability

This command is available on the [PoE](#) devices listed in [Extreme Networks PoE Devices](#).

show inline-power fast ports

```
show inline-power fast ports {port_list}
```

Description

Displays fast [PoE](#) information by port.

Syntax Description

fast	Specifies showing fast PoE power status, which is the ability to provide power to devices at the time of switch power on without waiting for boot up based on last saved PoE state.
ports	Designates showing fast PoE information for specified ports.
<i>port_list</i>	Specifies on which ports to show fast PoE information.

Default

All ports are fast PoE disabled by default.

Usage Guidelines

This command shows the port level fast PoE configurations saved in hardware, along with the ones currently configured to allow you to see any mismatches. If the default configuration file has never been saved in this boot session, the hardware configurations shown are the initial running configurations from hardware at boot time. If the default configuration file has been saved in this boot session (by `save configuration`), the hardware configurations shown are the last configurations saved to hardware.

Example

The following example shows fast PoE information on ports 1:1-2, 2:1-5, 3:1-3:

# show inline-power fast ports 1:1-2,2:1-5,3:1-3						
Port	State	Enable	Operator Limit	Priority	Detection	
----	----	-----	-----	-----	-----	
1:1	HW	Enabled	28000 mW	Low	802.3af-only	
	Cfg	Enabled	28000 mW	Low	802.3af-only	
1:2	HW	Enabled	30000 mW	High	4-point legacy-and-802.3af	
	Cfg	Enabled	30000 mW	High	4-point legacy-and-802.3af	
3:1	HW	Enabled	30000 mW	Critical	802.3af-only	
	Cfg	Enabled	30000 mW	Critical	802.3af-only	
3:2	HW	Disabled	15000 mW	Low	802.3af-only	
	Cfg	Disabled	15000 mW	Low	802.3af-only	
3:3	HW	Enabled	30000 mW	High	802.3af-only	
	Cfg	Enabled	30000 mW	Low	802.3af-only	

History

This command was first available in ExtremeXOS 30.3.

Platform Availability

This command is available on the [PoE](#) devices listed in [Extreme Networks PoE Devices](#).

show inline-power info ports

```
show inline-power info {detail} ports {port_list}
```

Description

Displays inline power information for the specified ports.

Syntax Description

<i>port_list</i>	Specifies one or more ports. If you do not specify ports, information for all ports appears.
------------------	--

Default

N/A.

Usage Guidelines



Note

Ports in the denied or faulted state periodically display the searching state as the hardware retests the PD state.

You can use this command to generate a summary report or a detailed report.

Summary output displays the following inline power information for the specified ports:

- State—Displays the port power state:
 - Disabled.
 - Searching.
 - Delivering.
 - Faulted.
 - Disconnected.
 - Other.
 - Denied.
- PD's power class—Displays the class type of the connected PD:
 - "-----": disabled or searching.
 - "class0": class 0 device.
 - "class1": class 1 device.
 - "class2": class 2 device.
 - "class3": class 3 device.
 - "class4": class 4 device.
 - "class5": class 5 device.
 - "class6": class 6 device.
 - "class7": class 7 device.
 - "class8": class 8 device.
- Volts—Displays the measured voltage. A value from 0 to 2 is valid for ports that are in a searching state.
- Curr—Displays the measured current, in milliamperes, drawn by the PD.
- Power—Displays the measured power, in watts, supplied to the PD.
- Fault—Displays the fault value:
 - None.
 - UV/OV fault.
 - UV/OV spike.
 - Over current.
 - Overload.
 - Undefined.
 - Underload.
 - HW fault.
 - Discovery resistance fail.
 - Operator limit violation.
 - Disconnect.

- Discovery resistance, A2D failure.
- Classify, A2D failure.
- Sample, A2D failure.
- Device fault, A2D failure.
- Force on error.

The detail command lists all inline power information for the selected ports.

Detail output displays the following information:

- Configured Admin State—Displays the port's configured state; Enabled or Disabled.
- Inline Power State—Displays the port power state.
- MIB Detect Status—Displays the port state as reported by *SNMP*; valid values are as follows:
 - disabled
 - searching
 - delivering
 - fault
 - test
 - otherFault
 - denyLowPriority
- Label—Displays the port's configured label.
- Operator Limit—Displays the port's configured operator limit value.
- PD Class—Displays the class type of connected PD:
- Max Allowed Power—Displays the amount of maximum allowed power for a device of this class.
- Measured Power—Displays the measured power, in watts, supplied to the PD.
- Line Voltage—Displays the measured voltage. A value from 0 to 2 is valid for ports in a searching state.
- Current—Displays the measured current, in milliamperes, drawn by the PD.
- Fault Status—Displays the fault value.
- Detailed Status
- Priority
- Power Up Status—2-pair or 4-pair, 15W–90W
- Powered Channels—primary channel, alternative channel, or both
- Detection

Example

The following command displays summary inline power information for ports 1 to 3 on a switch:

```
show inline-power info ports 1-3
```

The following is sample output from this command:

Port (mA)	State (Watts)	Class	Volts	Curr	Power	Fault
1	delivering	class3	48.3	192	9.300	None
2	delivering	class3	48.3	192	9.300	None
3	searching	-----	0.0	0	0.0	None

The following command displays detail inline power information for port 1 on slot 3:

```
show inline-power info detail port 3:1
```

The following is sample output from this command:

```
Port 3:1
Configured Admin State: enabled
Inline Power State      : delivering
MIB Detect Status       : delivering
Label                   :
Operator Limit          : 16800 milliwatts
PD Class                 : class3
Max Allowed Power       : 15.400 W
Measured Power          : 9.400 W
Line Voltage            : 48.3 Volts
Current                 : 193 mA
Fault Status            : None
Detailed Status         :
```

The following command displays detail inline power information for port 27:

```
# show inline-power info detail ports 27
```

The following is sample output from this command:

```
Port 27

Configured Admin State: enabled
Inline Power State      : delivering
MIB Detect Status       : delivering
Label                   :
Operator Limit          : 90000 milliwatts
PD Class                 : class8
Max Allowed Power       : 90.0 W
Measured Power          : 33.300 W
Line Voltage            : 53.7 Volts
Current                 : 621 mA
Fault Status            : None
Detailed Status         : Delivering power to IEEE PD
Priority                 : low
Power Up Status         : 4-pair 60W
Powered Channels        : Primary and Alternative
Detection               : 802.3af-only
```

The following command displays detail inline power information with **class-based** configured for port 1:

```
# show inline-power info detail ports 1
```

The following is sample output for this command (lines 7 and 9):

```
Port 1

Configured Admin State: enabled
Inline Power State      : delivering
MIB Detect Status       : delivering
Label                   :
Operator Limit          : class-based (97000 milliwatts when not class-based)
PD Class                 : class6
Max Allowed Power       : 60.0 W
Measured Power          : 49.800 W
Line Voltage            : 54.2 Volts
Current                 : 919 mA
```

Fault Status	: None
Detailed Status	: Delivering power to IEEE PD
Priority	: low
Power Up Status	: 4-pair 90W
Powered Channels	: Primary and Alternative
Detection	: 802.3af-only

History

This command was first available in ExtremeXOS 11.1.

PD classes 7 and 8, and power up status and powered channels information for ExtremeSwitching X465 PoE switches with 60W ('u') and 90W ('w') ports was added in ExtremeXOS 30.2.

class-based information was added in ExtremeXOS 32.3.

Platform Availability

This command is available on the *PoE* devices listed in [Extreme Networks PoE Devices](#).

show inline-power slot

```
show inline-power [fast | perpetual] slot slot
```

Description

Displays inline power information for the specified node (slot) on SummitStacks.

Syntax Description

<i>slot</i>	Specifies the node (slot).
fast	Specifies showing fast PoE power status, which is the ability to provide power to devices at the time of switch power on without waiting for boot up based on last saved PoE state.
perpetual	Specifies showing perpetual PoE power status, which is the ability to preserve PoE power delivery to devices during reboot. Perpetual PoE is a switch-wide setting.

Default

N/A.

Usage Guidelines

- On SummitStacks, the output indicates the following inline power status for each system:
- Configured power:
 - Enabled.

- Disabled.
- System power surplus.
- Redundant power surplus.
- Power usage threshold.
- Disconnect precedence.
- Legacy—The status of the legacy mode, which allows detection of many non-standard PDs.

The output indicates the following inline power status information for each slot:

- Inline power status—The status of inline power. The status conditions are:
 - Enabled.
 - Disabled.
- Firmware status—The operational status of the slot. The status conditions are:
 - Operational.
 - Not operational.
 - Disabled.
 - Subsystem failure.
 - Card not present.
 - Slot disabled.
- Budgeted power—The amount of power, in watts, that is available to the slot.
- Measured power—The amount of power, in watts, that currently being used by the slot.

On a stack, if you do not specify a slot number, the command operates on all active nodes. This command operates only on nodes in the active topology.

Using the **fast** option of this command shows the slot-level fast PoE information. It shows the last fast PoE boot result. It also shows the slot-level PoE configurations saved in hardware along with the ones currently configured so that you can see any mismatches. If the default configuration file has never been saved in this boot session, the hardware configurations shown in this command are the initial running configurations from hardware at boot time. If the default configuration file has been saved in this boot session (by `save configuration`), the hardware configurations shown in this command are the last configurations saved to hardware.

Example

The following command displays inline power information for node (slot) 3 on a SummitStack:

```
# show inline-power slot 3
Inline Power System Information
Configured                : Enabled
System Power Surplus      : 1500 Watts available for budgeting
Redundant Power Surplus   : 465 Watts available for budgeting to maintain N+1
Power Usage Threshold     : 70 percent (per slot)
Disconnect Precedence     : lowest-priority
Legacy Mode               : Disabled
Budgeted      Measured
Slot  Inline-Power  Firmware Status  Power (Watts)  Power (Watts)
3     Enabled      Operational      50 W          9 W
4     Enabled      Card Not Present  ( 50 W)       n/a
7     Enabled      Operational      50 W          0 W
Note: A budget value in parentheses is not allocated from the system power
```

The following example shows fast PoE status on a stack:

```
# show inline-power fast
Slot  State  Enable  Disconnect  Last Fast Boot Status
----  -
1      HW      Enabled  Deny-port    Succeeded
      Cfg      Enabled  Deny-port
3      HW      Disabled Deny-port    Disabled
      Cfg      Disabled Deny-port
4      HW      Enabled  Lowest-priority  Mismatch
      Cfg      Enabled  Deny-port
5      HW      Enabled  Deny-port    Failed (Query failed)
      Cfg      Enabled  Deny-port
6      HW      Enabled  Deny-port    Failed (Firmware upgrade)
      Cfg      Enabled  Deny-port
7      HW      Enabled  Deny-port    Failed (In reset)
      Cfg      Enabled  Deny-port
8      HW      Enabled  Deny-port    Failed (Invalid config)
      Cfg      Enabled  Deny-port

State: (HW) Hardware state per-slot, (Cfg) Configured state per-slot.
Fast Boot Status:
  Succeeded - Fast PoE boot succeeded and configuration matches.
  Disabled  - Fast PoE was disabled in hardware at boot time.
  Mismatch  - Fast PoE booted but the configuration in hardware
              did not match the configuration file.
  Failed    - Fast PoE boot failed and PoE was reset.
```

The following example shows perpetual PoE status on slot 3:

```
# show inline-power perpetual on slot 3
Perpetual POE is Enabled
```

History

This command was first available in ExtremeXOS 11.1.

The **fast** and **perpetual** PoE options were added in ExtremeXOS 30.3.

Platform Availability

It is available on SummitStack when the stack contains switches listed in [Extreme Networks PoE Devices](#).

show inline-power stats ports

```
show inline-power stats ports {port_list}
```

Description

Displays inline power statistics for the specified ports.

Syntax Description

<i>port_list</i>	Specifies one or more slots and ports. If you do not specify ports, information for all ports appears.
------------------	--

Default

N/A.

Usage Guidelines

The output displays the following inline power statistics for the specified ports:

- State—Displays the port power state:
 - Disabled.
 - Searching.
 - Delivering.
 - Faulted.
 - Disconnected.
 - Other.
 - Denied.
- PD's power class—Displays the class type of the connected PD:
 - "-----": disabled or searching.
 - "class0": class 0 device.
 - "class1": class 1 device.
 - "class2": class 2 device.
 - "class3": class 3 device.
 - "class4": class 4 device.
- Absent—Displays the number of times the port was disconnected.
- InvSig—Displays the number of times the port had an invalid signature.
- Denied—Displays the number of times the port was denied.
- Over-current—Displays the number of times the port entered an overcurrent state.
- Short—Displays the number of times the port entered undercurrent state.

Example

The following command displays inline power configuration information for ports 1 to 10 in node (slot) 3 on a SummitStack:

```
show inline-power stats ports 3:1-10
```

Following is sample output from this command:

STATISTICS		COUNTERS					
Port	State	Class	Absent	InvSig	Denied	OverCurrent	Short
3:1	delivering	class3	0	0	0	18	0
3:2	delivering	class3	0	0	0	0	0
3:3	searching	class0	0	0	0	0	0
3:4	searching	class0	0	0	0	0	0
3:5	searching	class0	0	0	0	0	0
3:6	searching	class0	0	0	0	0	0
3:7	searching	class0	0	0	0	0	0
3:8	searching	class0	0	0	0	0	0

3:9	searching	class0	0	0	0	0	0
3:10	searching	class0	0	0	0	0	0

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on the *PoE* devices listed in [Extreme Networks PoE Devices](#).

show inline-power stats slot

```
show inline-power stats slot slot
```

Description

Displays inline power statistics for the specified slot on SummitStacks.

Syntax Description

<i>slot</i>	Specifies the slot.
-------------	---------------------

Default

N/A.

Usage Guidelines

Use this command to produce a report that shows the firmware status and version plus how many ports are currently faulted, powered, and waiting for power for the selected slots. Unlike the values displayed with the `show inline-power stats ports` command, these values (displayed with the `show inline-power stats slot` command) are current readings, not cumulative counters.

On a stack, if you do not specify a slot number, the command operates on all active nodes. This command operates only on nodes in the active topology.

Example

The following command displays inline power statistics information for node (slot) 3 on a SummitStack:

```
show inline-power stats slot 3
```

Following is sample output from this command:

```
Inline-Power Slot Statistics
Slot: 3
Firmware status           : Operational
Firmware revision         : 292b1
```

```
Total ports powered      : 7
Total ports awaiting power : 41
Total ports faulted      : 0
Total ports disabled     : 0
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on SummitStacks when the stack contains switches listed in [Extreme Networks PoE Devices](#).

show inline-power stats

```
show inline-power stats
```

Description

Displays inline power statistics for the specified switch.

Syntax Description

There are no variables or parameters for this command.

Default

N/A.

Usage Guidelines

Use this command to produce a report that shows the firmware status and version plus how many ports are currently faulted, powered, and waiting for power for the switch. Unlike the values displayed with the `show inline-power stats ports` command, these values are current readings, not cumulative counters.

Example

The following command displays inline power statistics information:

```
show inline-power stats
```

Following is sample output from this command:

```
Inline-Power Slot Statistics
Firmware status      : Operational
Firmware revision    : 292b1
Total ports powered  : 7
Total ports awaiting power : 17
```



```
Total ports faulted      : 0
Total ports disabled     : 0
```

History

This command was first available in ExtremeXOS 11.5.

Platform Availability

This command is available on the switches listed in [Extreme Networks PoE Devices](#).

show inline-power

```
show inline-power [fast | perpetual]
```

Description

Displays inline power status information for the specified *PoE* switch.

Syntax Description

fast	Specifies showing fast PoE power status, which is the ability to provide power to devices at the time of switch power on without waiting for boot up based on last saved PoE state.
perpetual	Specifies showing perpetual PoE power status, which is the ability to preserve PoE power delivery to devices during reboot. Perpetual PoE is a switch-wide setting.

Default

N/A.

Usage Guidelines

The output varies depending on the PoE device you are using.



Note

For additional information on inline power parameters, see the [show power budget](#) command.

Example

The following command displays inline power status for the switch:

```
# show inline-power
Inline Power System Information
Configured                : Enabled
Power Usage Threshold     : 70 percent
```

Firmware Status	Power (Watts)	Power (Watts)	Legacy
Operational	405 W	0 W	Disabled

The following example displays fast PoE status:

```
# show inline-power fast
State   Enable   Disconnect   Last Fast Boot Status
-----
HW      Disabled Deny-port    Disabled
Cfg     Disabled Deny-port

State: (HW) Hardware state per-slot, (Cfg) Configured state per-slot.
Fast Boot Status:
  Succeeded - Fast PoE boot succeeded and configuration matches.
  Disabled  - Fast PoE was disabled in hardware at boot time.
  Mismatch  - Fast PoE booted but the configuration in hardware
              did not match the configuration file.
  Failed    - Fast PoE boot failed and PoE was reset.
```

The following example displays perpetual PoE status:

```
# show inline-power perpetual
Perpetual POE is Enabled
```

The following example displays power budget information (line 7):

```
# show inline-power

Inline Power System Information
Configured           : Enabled
Power Usage Threshold : 70 percent
Disconnect Precedence : deny-port
Power Budget         : Maximum

Firmware Status      Budgeted   Measured
Power (Watts)        Power (Watts)
Operational          490 W         5 W
```

History

This command was first available in ExtremeXOS 11.1.

The **fast** and **perpetual** PoE options were added in ExtremeXOS 30.3.

Platform Availability

This command is available on the switches listed in [Extreme Networks PoE Devices](#).

show ip nat

```
show ip nat
```

Description

Shows Network Address Translation (NAT) information.

Syntax Description

ip	Shows Internet Protocol (IP) information.
nat	Shows NAT information.

Default

N/A.

Usage Guidelines

NAT maps IP addresses from one address domain (typically a private IP address space) to an another address domain (typically a public Internet IP address space) to provide transparent routing to end hosts. This translation is accomplished transparently by having a NAT device translate the IP address and/or Layer 4 port of the packets.

To enable or disable IP NAT, use the commands `enable ip nat` and `disable ip nat`.

Example

The following example shows IP NAT information:

```
# show ip nat
NAT                               : Disabled
Number of NAT VLANs               : 0
Maximum number of NAT VLANs      : 4
Number of NAT rules               : 0
Number of source NAT rules        : 0
Number of destination NAPT rules  : 0
Number of configured NAPT rules   : 0
Number of dynamic NAPT rules      : 0
Maximum number of NAT rules       : 1024
Aging Time                        : 20 minutes
```

History

This command was first available in ExtremeXOS 31.2.

Platform Availability

This command is available on the switches that support the IP NAT feature. For information about which switches support this and other features, see the [Switch Engine 32.3 Feature License Requirements](#) document.

show ip nat rule

```
show ip nat rule {detail}
```

Description

Displays information about IP Network Address Translation (NAT) rules.

Syntax Description

ip	Specifies Internet Protocol (IP).
nat	Specifies NAT.
rule	Specifies viewing NAT rule information.
detail	Specifies showing detailed NAT rule information.

Default

N/A.

Usage Guidelines

Example

The following example shows basic information about all IP NAT rules configured on the switch:

```
# show ip nat rule
Rule Name          State   VR          IPAddr
NewIPAddr
Type              EgressVLAN      Protocol  L4Port
NewL4Port
=====
baseNapt           Enabled  VR-Default   10.1.1.100/24
112.123.234.200
NAPT              out
-----
srcTxlate          Enabled  VR-Default   10.20.30.40/32
50.1.1.100
Src-NAT           out
-----
natIncomplete      Disabled  -----
-----
Src-NAT           -----
-----
-----
dnat-http-server   Enabled  VR-Default   40.1.1.2/32
10.1.1.100
Dst-NAPT          out
UDP              8080
80
=====
=====
NAT                : Enabled
Total number of NAT rules      : 4
Total number of source NAT rules : 2
Total number of destination NAPT rules : 1
Total number of configured NAPT rules : 1
Total number of dynamic NAPT rules : 0
Maximum number of NAT rules    : 1024
Aging Time                : 10 minutes
```

The following example shows detailed information about all IP NAT rules configured on the switch:

```
show ip nat rule detail
Rule Name      : baseNapt
Type           : NAPT
Admin State    : Enabled
```

```

VR Name           : VR-Default
New VR Name       : VR-Default
Egress VLAN Name  : out
Source IP Address : 10.1.1.100
Source IP Mask    : 255.255.255.0
New Source IP Address : 112.123.234.200
Statistics Monitor : Off

Rule Name         : srcTxlate
Type              : Src-NAT
Admin State       : Enabled
VR Name           : VR-Default
New VR Name       : VR-Default
Egress VLAN Name  : out
Source IP Address : 10.20.30.40
Source IP Mask    : 255.255.255.255
New Source IP Address : 50.1.1.100
Statistics Monitor : Off

Rule Name         : natIncomplete
Type              : Src-NAT
Admin State       : Disabled
VR Name           :
New VR Name       :
Egress VLAN Name  :
Source IP Address :
Source IP Mask    :
New Source IP Address :
Statistics Monitor : Off

Rule Name         : dnat-http-server
Type              : Dst-NAPT
Admin State       : Enabled
VR Name           : VR-Default
New VR Name       : VR-Default
Egress VLAN Name  : out
Dest IP Address   : 40.1.1.2
Dest IP Mask      : 255.255.255.255
New Dest IP Address : 10.1.1.100
Protocol          : UDP (17)
Dest Port         : 8080
New Dest Port     : 80
Statistics Monitor : Off

=====
NAT                : Enabled
Total number of NAT rules      : 4
Total number of source NAT rules : 2
Total number of destination NAPT rules : 1
Total number of configured NAPT rules : 1
Total number of dynamic NAPT rules : 0
Maximum number of NAT rules    : 1024
Aging Time                  : 20 minutes

```

History

This command was first available in ExtremeXOS 31.2.

Platform Availability

This command is available on the switches that support the IP NAT feature. For information about which switches support this and other features, see the [Switch Engine 32.3 Feature License Requirements](#) document.

show ip nat rule statistics

```
show ip nat rule {rule_name} statistics {no-refresh}
```

Description

Shows statistics IP Network Address Translation (NAT) statistics (byte/packet counters) in the outward direction.

Syntax Description

ip	Specifies Internet Protocol (IP).
nat	Specifies NAT.
rule	Specifies NAT rules.
<i>rule_name</i>	Specifies displaying statistics only for the specified NAT rule.
statistics	Specifies showing IP NAT rule statistics.
no-refresh	Specifies displaying page-by-page without auto-refresh.

Default

N/A.

Usage Guidelines

Example

The following example shows NAT statistics in the outward direction:

```
# show ip nat rule statistics
NAT Rule Statistics                               Sat Sep 11 16:27:23 2019
NAT Rule Name                                Tx Packet                                Tx Byte
=====
nat59                                         2168                                     138752
nat60                                         2168                                     138752
nat61                                         2168                                     138752
nat62                                         2168                                     138752
nat63                                         2168                                     138752
nat64                                         2168                                     138752
nat65                                         2168                                     138752
nat66                                         2168                                     138752
nat67                                         2168                                     138752
nat68                                         2168                                     138752
nat69                                         2168                                     138752
nat70                                         2168                                     138752
```

```

nat71          2168          138752
nat72          2168          138752
nat73          2168          138752
nat74          2168          138752
nat75          2168          138752
nat76          2168          138752
nat77          2168          138752
=====
0->Clear Counters  U->page up  D->page down  ESC->exit

```

History

This command was first available in ExtremeXOS 31.2.

Platform Availability

This command is available on the switches that support the IP NAT feature. For information about which switches support this and other features, see the [Switch Engine 32.3 Feature License Requirements](#) document.

show ip nat vlan

```
show ip nat vlan {vlan_name}
```

Description

Shows Network Address Translation (NAT) VLAN information.

Syntax Description

ip	Shows Internet Protocol (IP) information.
nat	Shows NAT information.
vlan	Specifies VLAN NAT information.
<i>vlan_name</i>	Specifies which VLAN to show NAT information for.

Default

N/A.

Usage Guidelines

To view information for a specific VLAN, specify the VLAN name when running this command.

Example

The following example shows VLAN NAT information:

```

# show ip nat vlan
VLAN Name          Direction
=====

```

```

engineering          Both
in                   Ingress
out                  Egress
=====
Total Number of NAT VLANs   : 3
Maximum number of NAT VLANs : 4

```

History

This command was first available in ExtremeXOS 31.2.

Platform Availability

This command is available on the switches that support the IP NAT feature. For information about which switches support this and other features, see the [Switch Engine 32.3 Feature License Requirements](#) document.

show ip nat vlan counters

```
show ip nat vlan counters {vlan_name}
```

Description

Shows the number of packets translated and discarded in the Network Address Translation (NAT) VLAN egress and ingress directions.

Syntax Description

ip	Shows Internet Protocol (IP) information.
nat	Shows NAT information.
vlan	Shows VLAN NAT information.
counters	Shows NAT VLAN counter information
<i>vlan_name</i>	Specifies which VLAN to show NAT information for.

Default

N/A.

Usage Guidelines

Source translations and discards constitute the packet counts in the egress direction. Destination translations and discards constitute the packet counts in the ingress direction.

To clear counters, run the command `clear ip nat counters vlan {vlan_name}`.

Example

The following example shows NAT VLAN counter information:

```
# show ip nat vlan counters
```

VLAN Name	Source Translations	Source Discards	Destination Translations	Destination Discards
out	203284	0	0	0
in	0	0	0	0

History

This command was first available in ExtremeXOS 31.2.

Platform Availability

This command is available on the switches that support the IP NAT feature. For information about which switches support this and other features, see the [Switch Engine 32.3 Feature License Requirements](#) document.

show iparp

```
show iparp {ip_addr | mac | [{vlan} vlan_name | vlan vlan_list] |
    permanent} {port port {vr vr_name}}
```

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

<i>ip_addr</i>	Specifies an IP address.
<i>mac</i>	Specifies a MAC address.
<i>vlan_name</i>	Specifies a VLAN name.
<i>vlan_list</i>	Specifies a VLAN list of IDs.
permanent	Specifies permanent entries.
port	Show entries for a specific port.
<i>port</i>	Show only entries for this port number.
<i>vr_name</i>	Specifies a VR or VRF.

Default

Show all entries, except for proxy entries.

Usage Guidelines

Displays the IP ARP table, including:

- IP address
- MAC address
- Aging timer value
- VLAN name, VLAN ID and port number
- Flags

If you do not specify a VR or VRF, the command applies to the current VR context.

The show output displays the following information:

ARP address check	Whether IP ARP checking is enabled or disabled. IP ARP checking verifies if the ARP request's source address is in the receiving interface's subnet.
ARP refresh	Whether ARP refresh is enabled or disabled. ARP refresh is performed when an ARP entry's age is three-fourths of the timeout value.
Dynamic entries	The number of dynamic (learned ARP) entries in the table.
Pending entries	The number of sent ARP requests that have not yet received a response.
Static entries	The number of configured (static ARP) entries in the table.
Timeout	Timeout value for a dynamic (learned) ARP entry.
ARP Global Settings	
Max Entries	Maximum number of ARP entries.
Max Pending Entries	Maximum number of pending ARP entries.
Max Proxy Entries	Maximum number of proxy ARP entries.

Example

The following example displays the IP ARP table for the current VR or VRF context:

```
# show iparp
VR      Destination      Mac                Age  Static  VLAN      VID  Port
VR-Default  10.10.10.6      00:04:96:1f:a5:71    8    NO    bluered   4092  1
VR-Default  10.128.32.1     00:01:30:ba:6a:a0    0    NO    Default   4095
VR-Default  10.128.32.2     00:01:03:1c:ae:b0    5    NO    Default   4095
VR-Default  10.128.32.4     00:d0:59:17:74:83    3    NO    Default   4095
VR-Default  10.128.32.5     00:02:a5:c2:5c:dd    0    NO    Default   4095
VR-Default  10.128.32.6     00:12:3f:1c:f8:fb    5    NO    Default   4095
VR-Default  10.128.32.7     00:11:11:80:9c:b9    7    NO    Default   4095
VR-Default  10.128.32.8     00:11:43:53:8e:f1    0    NO    Default   4095
VR-Default  10.128.32.9     00:02:a5:bf:ac:70    7    NO    Default   4095
VR-Default  10.128.32.10    00:11:43:44:18:68   10    NO    Default   4095
VR-Default  10.128.32.11    00:12:3f:1c:e9:f2    0    NO    Default   4095
VR-Default  10.128.32.12    00:02:a5:bf:af:79    8    NO    Default   4095
VR-Default  10.128.32.13    00:11:43:40:89:91    0    NO    Default   4095
VR-Default  10.128.32.16    00:0f:1f:c9:2d:80    2    NO    Default   4095
VR-Default  10.128.32.17    00:06:5b:b1:6a:91    1    NO    Default   4095
VR-Default  10.128.32.19    00:11:43:3a:96:1d   10    NO    Default   4095
VR-Default  10.128.32.20    00:08:02:d5:c5:b7    6    NO    Default   4095
VR-Default  10.128.32.24    00:12:3f:0a:44:92   14    NO    Default   4095
```

```

VR-Default      10.128.32.26      00:50:04:ad:36:5e      6      NO      Default      4095
VR-Default      10.128.32.30      00:b0:d0:23:f2:9a      11     NO      Default      4095
VR-Default      10.128.32.54      00:b0:d0:59:e4:e2      6      NO      Default      4095
VR-Default      10.128.32.55      00:a0:c9:0c:41:de      3      NO      Default      4095
VR-Default      10.128.32.59      00:b0:d0:7c:d6:07      14     NO      Default      4095
VR-Default      10.128.32.99      00:04:96:05:00:03      13     NO      Default      4095
VR-Default      10.128.32.101     00:04:96:1f:a8:48      0      NO      Default      4095
VR-Default      10.128.32.104     00:30:48:41:ed:45      0      NO      Default      4095
VR-Default      10.128.32.105     00:30:48:41:ed:97      0      NO      Default      4095
VR-Default      10.128.32.106     00:01:30:23:c1:00      0      NO      Default      4095
VR-Default      10.128.32.108     00:04:96:1f:a5:71      0      NO      Default      4095
VR-Default      10.128.32.116     00:04:96:1f:a4:0e      0      NO      Default      4095
Dynamic Entries :          1          Static Entries      :          0
Pending Entries :          0
ARP address check:      Enabled          ARP refresh          :      Enabled
Timeout              :      20 minutes      ARP Sender-Mac Learning :      Disabled
Locktime             :      1000 milliseconds
Retransmit Time      :      1000 milliseconds
Reachable Time       :      900000 milliseconds (Auto)
Fast Convergence     :      Off

ARP Global Settings
Max Entries          :      4096
Max Pending Entries  :      256
Max Proxy Entries    :      256

```

History

This command was first available in ExtremeXOS 10.1.

The **vr** option was added in ExtremeXOS 11.0.

The `vlan_list` variable was added in ExtremeXOS 16.1.

For ARP, virtual network remote endpoints appear in the port column in ExtremeXOS 22.3

In Request, Out Request, Failed Requests, Proxy Answered, Rx Error, Rejected Count, Rejected Port, Max ARP entries, In Response, Out Response, Dup IP Addr, Rejected IP, Rejected I/F, and Max ARP pending entries output removed; ARP Global Settings Max Entries, Max Pending Entries, Max Proxy Entries output added in ExtremeXOS 30.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show iparp proxy

```
show iparp proxy {[ipNetmask | ip_addr mask]} {vr vr_name}
```

Description

Displays the proxy ARP table.

Syntax Description

<i>ipNetmask</i>	Specifies an IP address/mask length.
<i>ip_addr</i>	Specifies an IP address.
<i>mask</i>	Specifies a subnet mask.
<i>vr_name</i>	Specifies a VR or VRF.

Default

N/A.

Usage Guidelines

If no argument is specified, then all proxy ARP entries are displayed.

If you do not specify a VR or VRF, the command applies to the current VR context.

Example

The following command displays the proxy ARP table:

```
# show iparp proxy
VR          Destination      Mask           Mac              Refcnt  Always  Dynamic
VR-Default  1.1.0.4                   255.255.255.255 00:04:96:97:f5:32    0       NO      YES
VR-Default  10.1.1.200                255.255.255.255 vrrp              0       YES     YES
VR-Default  10.1.1.201                255.255.255.255 vrrp              0       NO      NO
VR-Default  10.1.1.202                255.255.255.255 00:11:22:33:44:55    0       NO      NO
VR-Default  10.1.1.204                255.255.255.255 vrrp              0       YES     YES
```

History

This command was first available in ExtremeXOS 10.1.

The **vr** option was added in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show iparp security

```
show iparp security [ {vlan} vlan_name | vlan vlan_list]
```

Description

Displays the IP ARP security violation information for one or all VLANs.

Syntax Description

<i>vlan_name</i>	Specifies a single VLAN for which to display security violation information.
<i>vlan_list</i>	Specifies a VLAN list of IDs to display security violation information.

Default

Shows security violation information for all VLANs except Mgmt.

Usage Guidelines

None.

Example

The following example displays IP ARP security violation information for all VLANs:

```
Switch.4 # show iparp security
Most Recent Violation
=====
Vlan          Security  Violations  Type   IP address      MAC          Port
=====
Default      ----
test         ----
Security Setting: (G) Gratuitous ARP Protection
Violation Type   : (g) Gratuitous ARP Violation
```

The following example displays IP ARP security violation information for VLAN Default:

```
Switch.5 # show iparp security "Default"
Most Recent Violation
=====
Vlan          Security  Violations  Type   IP address      MAC          Port
=====
Default      ----
Security Setting: (G) Gratuitous ARP Protection
Violation Type   : (g) Gratuitous ARP Violation
```

History

This command was first available in ExtremeXOS 10.1.

The **vr** option was added in ExtremeXOS 11.0.

The *vlan_list* variable was added in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show iparp stats

```
show iparp stats [[vr_name | vr {all | vr_name} ] { no-refresh | refresh}
                 {vr summary}

show iparp stats [vlan {all {vr vr_name}} | [ {vlan} vlan_name | vlan
                 vlan_list] {no-refresh | refresh}

show iparp stats ports {all | port_list} { no-refresh | refresh}
```

Description

Displays the IP ARP statistics for one or more VRs, VLANs, or ports.

Syntax Description

<i>vr_name</i>	Specifies a VR or VRF for which to display statistics.
<i>vlan_name</i>	Specifies a <u>VLAN</u> for which to display statistics.
<i>vlan_list</i>	Specifies a VLAN list of IDs to display statistics.
refresh	Continuous refresh of output.
<i>port_list</i>	Specifies a list of ports for which to display statistics.

Default

Shows all VLAN ARP statistics in a dynamic display.

Usage Guidelines

VLAN statistics and totals are displayed for a single VR. When you display IPARP statistics for one or all VLANs, the display includes the specified VLANs for the specified VR. If you do not specify a VR for a VLAN report, the display includes the specified VLANs for the current VR context.

Periodically as part of cleanup, failed entries go down and hence the total entries goes down. In certain scenarios, they may help detect a problem (memory leak or an attack). The `show iparp stats` command shows totals across all VRs, including VR-Mgmt.

Example

The following example displays ARP table statistics for all VRs and VRFs:

```
# show iparp stats vr all
IP ARP VR Statistics           Wed Apr 07 15:30:49 2010
ARP Total    Dynamic    Static    Pending
=====
VR-Default
96           89         5         0
VR-Mgmt
4            2         2         0
VR-SV_PPPOE
287          286         1         0
VR-NV_PPPOE
```

```

19          19          0          0
chicago2
50          44          5          0
Total for all VRs
456         440         13         0
=====
U->page up  D->page down  ESC->exit

```

The following example displays ARP table statistics for all VLANs in the current VR context:

```

# show iparp stats vlan all
IP ARP VLAN Statistics                               Wed Apr 07 15:30:49 2010
VLAN                      ARP Total                Dynamic                Static
=====
VLAN_06-AAR                94                89                5
VLAN_07-AAR               122               121               1
VLAN_02-BOT                43                42               1
=====
Totals for VR U3c-South.
Dynamic :          440    Static   :          13    Pending   :          0
=====
U->page up  D->page down  ESC->exit

```

The following example displays ARP table statistics for ports 1:1 to 1:17:

```

# show iparp stats ports 1:1-1:17
IP ARP Port Statistics                               Wed Apr 07 15:30:49 2010
Port                      Link State              ARP Total                Dynamic                Static
=====
1:1                        A                94                89                5
1:2                        A                37                37                0
1:3                        A               122               121               1
1:4                        R                 0                 0                0
1:5                        R                 0                 0                0
1:6                        A                43                43                0
1:7                        A               118               118                0
1:8                        R                 0                 0                0
1:9                        R                 0                 0                0
1:10                       A                 8                 8                0
1:11                       A                 8                 6                2
1:12                       A                41                41                0
1:13                       A                17                17                0
1:14                       R                 0                 0                0
1:15                       R                 0                 0                0
1:16                       A                 8                 8                0
1:17                       A                 8                 6                2
=====
Link State: A-Active, R-Ready, NP-Port Not Present, L-Loopback
U->page up  D->page down  ESC->exit

```

History

This command was first available in ExtremeXOS 12.4.2.

The **no-refresh** keyword was removed in ExtremeXOS 16.1.

The `vlan_list` variable was added in ExtremeXOS 16.1.

The **refresh** keyword was added in ExtremeXOS 16.1.

The ARP unneeded, failed, rejected, total entries, and last rejected information was removed in ExtremeXOS 30.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show ipconfig

```
show ipconfig {ipv4} {vlan vlan_name | tunnel tunnel_name}
```

Description

Displays configuration information for one or more VLANs in the current VR context.

Syntax Description

<i>vlan_name</i>	Specifies a VLAN name.
<i>tunnel_name</i>	Specifies an IPv6 tunnel.

Default

N/A.

Usage Guidelines

If no VLAN information is specified, then global IP configuration is displayed. Otherwise, specific VLAN information is displayed.

Example

```
# show ipconfig
  Use Redirects : Disabled
  IpOption LSRR : Enabled
  IpOption SSRR : Enabled
    IpOption RR : Enabled
    IpOption TS : Enabled
    IpOption RA : Enabled
  Route Sharing : Disabled
  Originated Packets : Don't require ipforwarding
  IP Fwding into LSP : Disabled
  Max Shared Gateways : Current: 4   Configured: 4
  Route Sharing Hash : XOR, Custom Method
  Ping Protection : Enabled for static routes Interval: 2 Misses: 3
  IP Anycast MAC : 00:00:AB:BA:BA:BA
  IRDP:
    Advertisement Address: 255.255.255.255
    Maximum Interval: 600
    Minimum Interval: 450      Lifetime: 1800      Preference: 0
  Interface      IP Address      Flags              nSIA
  data2          200.0.0.1        /24 EUf---R----- 0
```


inet	1.1.1.2	/24	EUf---R-----	0
mytun	2.0.0.2	/24	EU---R-----	0

The following example displays configuration information about tunnel "mytunnel" with TCP MSS adjustment configuration (lines 3 and 8) and MTU value of 1400 (line 3):

```
# show ipconfig mytunnel
Router Interface on mytunnel is enabled and up.
  inet 2.0.0.1/24  broadcast 2.0.0.255  Mtu 1500  TCP-MSS  1360

Flags:
  BOOTP Host  NO   DirBcstHwFwd  NO           Fwd Bcast  NO   IgnoreBcast  YES
  IP Fwding   YES   IPmc Fwd   NO   Multinetted VLAN  NO   IRDP Advert  NO
  Send Redir  YES           VRRP      NO           Unicast RPF  NO  TCP Adjust MSS  ON
```

History

This command was first available in ExtremeXOS 10.1.

The **ipv4** keyword was added in ExtremeXOS 11.2.

This command changed to display information for the current VR context in ExtremeXOS 12.5.

Ping protection information added in ExtremeXOS 22.1.

Route sharing hashing algorithm (default or custom) added in ExtremeXOS 22.1.

The D (duplicate address detected on VLAN), t (tentative address), A (address mask reply enabled), M (send parameter problem enabled), P (send port unreachable enabled), T (time stamp reply enabled), u (send unreachable enabled), and X (send time exceeded enabled) flags were removed in ExtremeXOS 30.1.

IP anycast information was added in ExtremeXOS 30.6.

TCP MSS adjustment and IP MTU information were added in ExtremeXOS 31.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show ipconfig ipv6

```
show ipconfig ipv6 {vlan vlan_name | tunnel tunnel_name}
```

Description

Displays configuration information for one or more interfaces in the current VR context.

Syntax Description

<i>vlan_name</i>	Specifies an IPv6 configured VLAN .
<i>tunnel_name</i>	Specifies an IPv6 tunnel.

Default

N/A.

Usage Guidelines

If no interface is specified, then global IPv6 configuration is displayed. Otherwise, specific interface(s) will be displayed. Global IPv6 configuration information includes:

- IPv6 address/netmask/etc.
- IPv6 forwarding information/IPv6 multicast forwarding information

Example

The following example displays configuration information on a VLAN "accounting":

```
show ipconfig ipv6 vlan accounting
```

The current and configured values for **max-gateways** now apply to IPv6 gateway sets as well as IPv4, so these values are added to the output:

```
# show ipconfig ipv6
Route Sharing : Disabled
ICMP Redirect for Fast Path : Disabled
Max Shared Gateways : Current: 32 Configured: 32
```

The configuration settings for static route ping protection enable/disable, interval, and misses also appear:

```
# show ipconfig ipv6
Route Sharing : Enabled
Route Compression : Enabled
ICMP Redirect for Fast Path : Disabled
Max Shared Gateways : Current: 16 Configured: 16
Route Sharing Hash : CRC Lower
Ping Protection : Enabled for static routes Interval: 2 Misses: 3
IPv6 Anycast MAC : 00:00:AB:BA:BA:BA
```

The following example displays configuration information for ICMP multicast and anycast echo requests (lines 4-5):

```
# show ipconfig ipv6
Route Sharing : Disabled
Route Compression : Enabled
Ignore ICMP Multicasts : Enabled
Ignore ICMP Anycasts : Enabled
ICMP Redirect for Fast Path : Disabled
Max Shared Gateways : Current: 16 Configured: 16
Route Sharing Hash : CRC Lower, Default Method
Ping Protection : Enabled for static routes Interval: 2 Misses: 3
IPv6 Anycast MAC : None
```

The following example displays configuration information with TCP MSS adjustment support (lines 2 and 9):

```
# show ipconfig ipv6 v61
Router Interface on v61 is enabled. MTU: 1500 TCP-MSS: 0
IPv6 Hop Limit: 64
Locally registered unicast addresses:
```

```

    fe80::a01:101%v61/64 (Tentative)
  Flags:
    IPv6 Forwarding: YES    Accept recvd RA:  NO
    Send redirects: YES    Accept redirects:  NO
    TCP Adjust MSS:  OFF

```

History

This command was first available in ExtremeXOS 11.2.

This command changed to display information for the current VR context in ExtremeXOS 12.5.

The show output for **max-gateways** was added in ExtremeXOS 15.3.

Ping protection information added in ExtremeXOS 22.1.

Route sharing hashing algorithm (default or custom) added in ExtremeXOS 22.1.

IP anycast information was added in ExtremeXOS 30.6.

Enabling or disabling replies to multicast and anycast echo requests was added in ExtremeXOS 31.3.

TCP MSS adjustment information was added in ExtremeXOS 31.6.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

show ipmroute

```
show ipmroute {source-net mask-len | source-net mask | summary} {vr vr-name}
```

Description

Displays the contents of the IP multicast routing table or the route origin priority.

Syntax Description

<i>source-net</i>	Specifies an IP address/mask length.
<i>mask-len</i>	Mask length for the IP multicast source's subnet. Range is 1-32.
<i>mask</i>	Specifies a subnet mask.
summary	Displays the statistics of multicast static routes.
<i>vr-name</i>	Specifies the virtual router to which the route is added.

Default

vr-name is the VR of the current CLI context.

Usage Guidelines

This command allows you to view the configured multicast static routes. You can specify the filtering criteria on this command to view only the desired route. The multicast static routes are displayed in ascending order of their prefix (same order as `show iproute displays`).

Example

The following example shows a multicast static route from a default virtual router:

```
# show ipmroute
Destination      Gateway      Mtr      Flags      Protocol      VLAN
Default Route    20.20.20.1   255      UG         None          pc4-1
*1.1.0.0/16      20.20.20.1   10       UG         bgp           pc4-1
*11.0.0.0/8      30.30.30.1   12       U-         None          pc5-3
11.22.0.0/16     20.20.20.1   10       UG         None          pc4-1
*11.22.33.0/24   30.30.30.1   8        U-         None          pc5-3
11.22.33.44/32   20.20.20.1   4        UG         None          pc4-1
*12.0.0.0/8      20.20.20.1   0        UG         None          pc4-1
12.24.0.0/16     30.30.30.1   0        U-         None          pc5-3
*12.24.48.96/32  30.30.30.1   2        U-         ospf-externl  pc5-3
44.66.0.0/16     30.30.30.1   0        U-         None          pc5-3
Flags: (*) Active, (G) Gateway, (U) Up
Mask distribution:
1 default routes 2 routes at length 8
4 routes at length 16 1 routes at length 24
2 routes at length 32
Total number of multicast static routes = 10
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IPv4 multicast feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

show iproute

```
show iproute {ipv4} {priority | vlan vlan_name | permanent | ip_address
netmask | summary} {multicast | unicast} {vr vrname}}
```

Description

Displays the contents of the IP routing table or the route origin priority.

Syntax Description

priority	Displays the priority values for each route origin type.
<i>vlan_name</i>	Specifies a <u>VLAN</u> name.

permanent	Specifies permanent routing.
<i>ip_address</i>	Specifies an IP address.
<i>netmask</i>	Specifies a subnet mask.
summary	Displays summary information.
multicast	Displays information for IPv4 multicast routes only.
unicast	Displays information for IPv4 unicast routes only.
<i>vrname</i>	Specifies a VR or VRF.

Default

N/A.

Usage Guidelines

A c flag in the **Flags** column indicates a compressed route resulting from enabling compression using the enable iproute compression command. The total number of compressed routes is also shown.

All routes that are provided to the FIB display the f flag.

If you do not specify a VR or VRF, the command applies to the current VR context.

Example

The following example displays detailed information about all IP routing:

```
# show iproute vr vr-b
Ori Destination Gateway Mtr Flags VLAN Duration
#d 1.1.1.100/32 1.1.1.100 1 U-----um--f- l3vni_vr_b 0d:0h:1m:17s
*evn 2.2.2.2/32 2.2.2.2 1 UGHD---u---f- l3vni_vr_b 0d:0h:0m:8s
*evn 2.2.2.200/32 2.2.2.2 1 UGHD---u---f- l3vni_vr_b 0d:0h:0m:7s
#d 20.1.101.0/24 20.1.101.1 1 U-----um--f- tenant_vlan_101 0d:0h:1m:17s
*evn 20.1.102.0/24 2.2.2.2 1 UG-D---u---f- l3vni_vr_b 0d:0h:0m:7s
*evn 20.1.102.2/32 2.2.2.2 1 UGHD---u---f- l3vni_vr_b 0d:0h:0m:7s
*evn 99.0.0.0/8 2.2.2.2 1 UG-D---u---f- l3vni_vr_b 0d:0h:0m:7s

Origin(Ori): (ap) Auto-peering, (b) BlackHole, (be) EBGp, (bg) BGP, (bi) IBGP,
              (bo) BOOTP, (ct) CBT, (d) Direct, (df) DownIF, (dv) DVMRP,
              (e1) ISISL1Ext, (e2) ISISL2Ext, (evn) EVPN, (h) Hardcoded,
              (hm) Host-mobility, (i) ICMP, (i1) ISISL1 (i2) ISISL2, (is) ISIS, (mb) MBGP,
              (mbe) MBGPExt, (mbi) MBGPInter, (mp) MPLS Lsp,
              (mo) MOSPF, (o) OSPF, (o1) OSPFExt1, (o2) OSPFExt2, (oa) OSPFIntra
              (oe) OSPFAsExt, (or) OSPFInter, (pd) PIM-DM, (ps) PIM-SM,
              (r) RIP, (ra) RtAdvrt, (s) Static, (sv) SLB_VIP, (un) UnKnown,
              (*) Preferred unicast route (@) Preferred multicast route,
              (#) Preferred unicast and multicast route.

Flags: (b) BFD protection requested, (B) BlackHole, (c) Compressed, (D) Dynamic,
        (f) Provided to FIB, (G) Gateway, (H) Host Route,
        (I) ICMP ping protection requested, (l) Calculated LDP LSP,
        (L) Matching LDP LSP, (m) Multicast, (p) BFD protection active,
        (P) LPM-routing, (R) Modified, (s) Static LSP, (S) Static,
        (t) Calculated RSVP-TE LSP, (T) Matching RSVP-TE LSP, (u) Unicast,
        (U) Up, (3) L3VPN Route.
```

```

MPLS Label: (S) Bottom of Label Stack
Mask distribution:
    1 routes at length 8          2 routes at length 24
    4 routes at length 32

Route Origin distribution:
    2 routes from Direct          5 routes from EVPN

Total number of routes = 7
Total number of compressed routes = 0

```

The following example displays detailed priority values for VR 'vr-test':

```

# show iproute priority vr vr-test
Direct                10
MPLS                  20
Blackhole             50
Static                1100
HostMobility          1150
ICMP                  1200
EVPN                  1698
Autopeering           1699
EBGP                  1700
IBGP                  1900
OSPFIntra             2200
OSPFInter             2300
Isis                  2350
IsisL1                2360
IsisL2                2370
RIP                   2400
OSPFAsExt             3100
OSPFExt1              3200
OSPFExt2              3300
IsisL1Ext             3400
IsisL2Ext             3500
Bootp                 5000
OSPFInterVr           5010

```

The following example displays detailed information for VR 'vrf-test':

```

# show iproute vr vrf-test
Ori Destination      Gateway      Mtr  Flags      VLAN      Duration
#oiv 40.0.0.0/24      11.0.0.2    5010  UG-D---um--f- ce2_pe2_1  0d:0h:1m:15s

Origin(Ori): (ap) Auto-peering, (b) BlackHole, (be) EBGp, (bg) BGP, (bi) IBGP,
              (bo) BOOTP, (ct) CBT, (d) Direct, (df) DownIF, (dv) DVMRP,
              (e1) ISISL1Ext, (e2) ISISL2Ext, (evn) EVPN, (h) Hardcoded,
              (hm) Host-mobility, (i) ICMP, (i1) ISISL1 (i2) ISISL2, (is) ISIS, (mb) MBGP,
              (mbe) MBGPExt, (mbi) MBGPInter, (mp) MPLS Lsp,
              (mo) MOSPF (o) OSPF, (o1) OSPFExt1, (o2) OSPFExt2, (oa) OSPFIntra
              (oe) OSPFAsExt, (oiv) OSPFInterVr, (or) OSPFInter, (pd) PIM-DM, (ps) PIM-SM,
              (r) RIP, (ra) RtAdvrt, (s) Static, (sv) SLB_VIP, (un) UnKnown,
              (*) Preferred unicast route (@) Preferred multicast route,
              (#) Preferred unicast and multicast route.

```

History

This command was first available in ExtremeXOS 10.1.

The **ipv4** keyword was added in ExtremeXOS 11.2.

- The c flag was added in ExtremeXOS 12.0.
- The f flag was added in ExtremeXOS 12.2.2.
- The l flag showing *ICMP* ping protection was added in ExtremeXOS 22.1.
- The ap flag showing auto-peering information was added in ExtremeXOS 22.5.
- The evn origin prefix showing EVPN information was added in ExtremeXOS 30.7.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show iproute bfd

```
show iproute bfd
```

Description

Shows BFD protection gateway status for static routes added with BFD protection.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Example

The following example shows BFD protection gateway status for static routes added with BFD protection:

```
# show iproute bfd
Strict BFD for static routes : Enabled
```

Gateway	BFD Session State	# Routes	VR
1.1.1.2	Active (Down)	1	VR-Default
1.1.1.3	Pending		

History

- This command was first available in ExtremeXOS 12.4.
- Strict BFD session protection for static routes status was added in ExtremeXOS 30.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show iproute ipv6 origin

```
show iproute ipv6 origin [auto-peering | direct | static | blackhole
| ripng | ospfv3 | ospfv3-intra | ospfv3-inter | ospfv3-extern1 |
ospfv3-extern2 | isis | isis-level-1 | isis-level-2 | isis-level-1-
external | isis-level-2-external | bgp | ibgp | ebgp | bootp | host-
mobility] {vr vr_name}
```

Description

Displays the contents of the IPv6 routing table for routes with the specified origin.

Syntax Description

direct	Specifies direct routes.
auto-peering	Specifies BGP auto-peering routes.
static	Specifies static routes.
blackhole	Specifies blackhole routes.
ripng	Specifies <i>RIPng</i> routes.
ospfv3	Specifies <i>OSPFv3</i> routes.
ospfv3-intra	Specifies OSPFv3 Intra routing.
ospfv3-inter	Specifies OSPFv3 Inter routing.
ospfv3-extern1	Specifies OSPFv3 External 1 routing.
ospfv3-extern2	Specifies OSPFv3 External 2 routing.
isis	Specifies IS-IS routes.
isis-level-1	Specifies IS-IS Level 1 routing.
isis-level-2	Specifies IS-IS Level 2 routing.
isis-level-1-external	Specifies IS-IS Level 1 External routing.
isis-level-2-external	Specifies IS-IS Level 2 External routing.
bgp	Specifies Border Gateway Protocol.
ibgp	Specifies Interior <i>BGP</i> route.
ebgp	Specifies Exterior BGP route.
bootp	Specifies BOOTP route.
host-mobility	Specifies host-mobility routes.
<i>vr_name</i>	Specifies a VR or VRF.

Default

N/A.

Usage Guidelines

None.

Example

The following example displays the RIPng routes:

```
show iproute ipv6 origin ripng
```

History

This command was first available in ExtremeXOS 11.2.

The **bootp** attribute was added in 15.7.1

The **auto-peering** option was added in ExtremeXOS 22.5.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

show iproute ipv6

```
show iproute ipv6 {priority | vlan vlan_name | tunnel tunnel_name |  
                  ipv6Netmask | summary {multicast | unicast}} {vr vr_name}}
```

Description

Displays the contents of the IPv6 routing table.

Syntax Description

priority	Displays the priority values for each route origin type.
<i>vlan_name</i>	Specifies a <u>VLAN</u> name.
<i>tunnel_name</i>	Specifies a tunnel name.
<i>ipv6Netmask</i>	Specifies an IPv6 address/prefix length.
summary	Specifies summary information.
<i>vr_name</i>	Specifies a VR or VRF.

Default

N/A.

Usage Guidelines

If you do not specify a VR or VRF, the command applies to the current VR context.

Example

The following example displays detailed information about all IPv6 routing:

```
Switch.18 # show iproute ipv6
Ori Destination                                     Mtr  Flags          Duration
Gateway Interface
#d 2001:db8::/64                                     1    U-----um--f 0d:0h:5m:31s
2001:db8::52                                         ixia
#or 2001:db8:2:78::/64                               50   UG-D---um--f 0d:0h:0m:1s
fe80::200:40ff:feba:a38e                             ixia
#or 2001:db8:2:79::/64                               50   UG-D---um--f 0d:0h:0m:1s
fe80::200:40ff:feba:a38e                             ixia
#or 2001:db8:2:7a::/64                               50   UG-D---um--f 0d:0h:0m:1s
fe80::200:40ff:feba:a38e                             ixia
#or 2001:db8:2:7b::/64                               50   UG-D---um--f 0d:0h:0m:1s
fe80::200:40ff:feba:a38e                             ixia
#or 2001:db8:2:7c::/64                               50   UG-D---um--f 0d:0h:0m:1s
fe80::200:40ff:feba:a38e                             ixia
#or 2001:db8:2:7d::/64                               50   UG-D---um--f 0d:0h:0m:1s
fe80::200:40ff:feba:a38e                             ixia
#or 2001:db8:2:7e::/64                               50   UG-D---um--f 0d:0h:0m:1s
fe80::200:40ff:feba:a38e                             ixia
#or 2001:db8:2:7f::/64                               50   UG-D---um--f 0d:0h:0m:1s
fe80::200:40ff:feba:a38e                             ixia
#or 2001:db8:2:80::/64                               50   UG-D---um--f 0d:0h:0m:1s
fe80::200:40ff:feba:a38e                             ixia
#or 2001:db8:2:81::/64                               50   UG-D---um--f 0d:0h:0m:1s
fe80::200:40ff:feba:a38e                             ixia
#d fe80::%ixia/64                                    1    U-----um--f 0d:0h:5m:31s
fe80::204:96ff:fe27:8697                             ixia
Origin(Ori): (b) BlackHole, (be) EBGp, (bg) BGP, (bi) IBGP, (bo) BOOTP
(ct) CBT, (d) Direct, (df) DownIF, (dv) DVMRP, (el) ISISL1Ext
(e2) ISISL2Ext, (h) Hardcoded, (i) ICMP, (il) ISISL1 (i2) ISISL2
(is) ISIS, (mb) MBGP, (mbe) MBGPExt, (mbi) MBGPInter, (ma) MPLSIntra
(mr) MPLSInter, (mo) MOSPF (o) OSPFv3, (ol) OSPFv3Ext1, (o2) OSPFv3Ext2
(oa) OSPFv3Intra, (oe) OSPFv3AsExt, (or) OSPFv3Inter, (pd) PIM-DM, (ps) PIM-SM
(r) RIPng, (ra) RtAdvrt, (s) Static, (sv) SLB_VIP, (un) UnKnown
(*) Preferred unicast route (@) Preferred multicast route
(#) Preferred unicast and multicast route
Flags: (B) BlackHole, (D) Dynamic, (G) Gateway, (H) Host Route
(I) ICMP ping protection requested,
(L) Matching LDP LSP, (l) Calculated LDP LSP, (m) Multicast
(P) LPM-
routing, (R) Modified, (S) Static, (s) Static LSP
(T) Matching RSVP-TE LSP (t) Calculated RSVP-TE LSP, (u) Unicast, (U) Up
(f) Provided to FIB (c) Compressed Route
Mask distribution:
12 routes at length 64
Route Origin distribution:
2 routes from Direct          10 routes from OSPFv3Inter
Total number of routes = 12
Total number of compressed routes = 0
```

The following example displays the IPv6 route origin priority:

```
Switch.4 # show iproute ipv6 priority
Direct          10
Blackhole       50
```

Static	1100
HostMobility	1150
ICMP	1200
Autopeering	1699
EBGP	1700
IBGP	1900
OSPFv3Intra	2200
OSPFv3Inter	2300
Isis	2350
IsisL1	2360
IsisL2	2370
RIPng	2400
OSPFv3AsExt	3100
OSPFv3Ext1	3200
OSPFv3Ext2	3300
IsisL1Ext	3400
IsisL2Ext	3500

History

This command was first available in ExtremeXOS 11.2.

The I flag showing ICMP ping protection was added in ExtremeXOS 22.1.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

show iproute mpls origin

```
show iproute mpls origin [bgp | blackhole | bootp | direct | ebgp | ibgp
| icmp | isis | isis-level-1 | isis-level-1-external | isis-level-2 |
isis-level-2-external | mpls evpn {signaling-protocol [ldp | rsvp-te
| static]} | ospf | ospf-extern1 | ospf-extern2 | ospf-inter | ospf-
intra | rip | static ] {unicast} {vr vrname}
```

Description

Displays the MPLS contents of the IP routing table for routes with the specified origin.

Syntax Description

bgp	Specifies <u>BGP</u> routes.
blackhole	Specifies blackhole routes.
bootp	Specifies BOOTP routes.
direct	Specifies direct routes.
ebgp	Specifies E-BGP routes.
ibgp	Specifies I-BGP routes.
icmp	Specifies <u>ICMP</u> routes.

isis	Specifies IS-IS routes.
isis-level-1	Specifies ISIS Level 1 routing.
isis-level-1-external	Specifies ISIS Level 1 External routing.
isis-level-2	Specifies ISIS Level 2 routing.
isis-level-2-external	Specifies ISIS Level 2 External routing.
mpls	Specifies MPLS routes. This option is available only on platforms that support the MPLS feature pack, which is described in the Switch Engine 32.3 Feature License Requirements document.
evpn	Specifies EVPN routes.
signaling-protocol [ldp rsvp-te static]	Specifies an MPLS signaling protocol. This option is available only on platforms that support the MPLS feature pack, which is described in the Switch Engine 32.3 Feature License Requirements document.
ospf	Specifies <i>OSPF</i> routes.
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 <i>RIP</i> routes.
static	Specifies static routes.
unicast	Displays unicast routes with the specified origin.
<i>vrname</i>	Specifies a VR or VRF.

Default

N/A.

Usage Guidelines

None.

Example

The following example displays all the MPLS routes that originate from BGP:

```
# show iproute mpls origin bgp
```

History

This command was first available in ExtremeXOS 12.2.2.

The **evpn** option was added in ExtremeXOS 30.7.

Platform Availability

This command is available on all platforms that support MPLS.

show iproute mpls

```
show iproute mpls {lsp lsp_name | vlan vlan_name | permanent |  
                  ip_address netmask | summary} {unicast} {vr vrname}
```

Description

Displays the *MPLS* contents of the IP routing table.

Syntax Description

<i>lsp_name</i>	Specifies an LSP name.
<i>vlan_name</i>	Specifies a <i>VLAN</i> name.
permanent	Specifies permanent routing.
<i>ip_address</i>	Specifies an IP address.
<i>netmask</i>	Specifies a subnet mask.
unicast	Displays unicast routes.
<i>vrname</i>	Specifies a VR or VRF.

Default

N/A.

Usage Guidelines

A c flag in the Flags column indicates a compressed route resulting from enabling compression using the enable iproute compression command. The total number of compressed routes is also shown.

All routes that are provided to the FIB display the f flag.

If you do not specify a VR or VRF, the command applies to the current VR context.

Example

The following example displays detailed information about all IP routing:

```
Switch.3 # show iproute mpls
```

History

This command was first available in ExtremeXOS 12.2.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show iproute multicast

```
show iproute {ipv4} [{vlan} name | [ipaddress netmask | ipNetmask] |
origin [direct | static | mbgp | imbgp | embgp]} multicast {vr
vr_name}
```

Description

Displays all or a filtered set of multicast routes in the IP multicast routing table.

Syntax Description

ipv4	Selects only IPv4 multicast routes.
<i>name</i>	Specifies a <u>VLAN</u> for which to display multicast routes.
<i>ipaddress netmask</i>	Specifies an IP address and network mask (in dotted decimal notation) for which to display multicast routes.
<i>ipNetmask</i>	Specifies the IP address and network mask in classless inter domain routing (CIDR) notation.
origin	Limits the displayed multicast routes to those generated by the specified origin. Origin options select direct routes, static routes, and routes created by the MBGP, IMBGP, and EMBGP protocols.
<i>vr_name</i>	Specifies the virtual router for which to display multicast routes.

Default

vr_name is the VR of the current CLI context.

Usage Guidelines

This command does not display unicast routes, which can be used for multicast traffic.

Example

The following example displays all the routes in multicast routing table:

```
# show iproute multicast
Ori Destination      Gateway      Mtr  Flags      VLAN      Duration
@d  3.3.3.3/32        3.3.3.3      1    U-----m--- lpbk      12d:1h:30m:36s
@d  28.0.0.0/24       28.0.0.15    1    U-----m--- trunk28    12d:1h:30m:36s
@mbe 77.0.0.0/24       50.1.10.21   1    UG---S--m--- toronto    0d:0h:41m:1s
@mbe 77.0.1.0/24       50.1.10.21   1    UG---S--m--- toronto    0d:0h:41m:1s
@mbe 77.0.2.0/24       50.1.10.21   1    UG---S--m--- toronto    0d:0h:41m:1s
@mbe 77.0.3.0/24       50.1.10.21   1    UG---S--m--- toronto    0d:0h:41m:1s
@mbe 77.0.4.0/24       50.1.10.21   1    UG---S--m--- toronto    0d:0h:41m:1s
@mbe 77.0.5.0/24       50.1.10.21   1    UG---S--m--- toronto    0d:0h:41m:1s
@mbe 77.0.6.0/24       50.1.10.21   1    UG---S--m--- toronto    0d:0h:41m:1s
@mbe 77.0.10.0/24      50.1.10.21   1    UG---S--m--- toronto    0d:0h:41m:1s
```

```

@ambe 77.0.11.0/24      50.1.10.21      1      UG---S--m--- toronto      0d:0h:41m:1s
@ambe 77.0.12.0/24      50.1.10.21      1      UG---S--m--- toronto      0d:0h:41m:1s
@ambe 77.0.13.0/24      50.1.10.21      1      UG---S--m--- toronto      0d:0h:41m:1s
@ambe 77.0.14.0/24      50.1.10.21      1      UG---S--m--- toronto      0d:0h:41m:1s
@d     82.0.0.0/24       82.0.0.15       1      U-----m--- trunk28-2    12d:1h:30m:36s
Origin(Ori): (b) BlackHole, (be) EBGp, (bg) BGP, (bi) IBGP, (bo) BOOTP
(ct) CBT, (d) Direct, (df) DownIF, (dv) DVMRP, (e1) ISISL1Ext
(e2) ISISL2Ext, (h) Hardcoded, (i) ICMP, (i1) ISISL1 (i2) ISISL2
(is) ISIS, (mb) MBGP, (mbe) MBGPExt, (mbi) MBGPInter, (mp) MPLS Lsp
(mo) MOSPF (o) OSPF, (o1) OSPFExt1, (o2) OSPFExt2
(oa) OSPFIntra, (oe) OSPFAsExt, (or) OSPFInter, (pd) PIM-DM, (ps) PIM-SM
(r) RIP, (ra) RtAdvrt, (s) Static, (sv) SLB_VIP, (un) UnKnown
(*) Preferred unicast route (@) Preferred multicast route
(#) Preferred unicast and multicast route
Flags: (B) BlackHole, (D) Dynamic, (G) Gateway, (H) Host Route
(L) Matching LDP LSP, (l) Calculated LDP LSP, (m) Multicast
(P) LPM-routing, (R) Modified, (S) Static, (s) Static LSP
(T) Matching RSVP-TE LSP, (t) Calculated RSVP-TE LSP, (u) Unicast, (U) Up
(f) Provided to FIB (c) Compressed Route
Mask distribution:
14 routes at length 24          1 routes at length 32
Route Origin distribution:
3 routes from Direct
Total number of routes = 15
Total number of compressed routes = 0

```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IPv4 multicast feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

show iproute origin

```

show iproute origin [auto-peering | bgp | blackhole | bootp | direct
| ebgp | embgp | ibgp | icmp | imbgp | isis | isis-level-1 | isis-
level-1-external | isis-level-2 | isis-level-2-external | mbgp | mpls
| evpn {signaling-protocol [ldp | rsvp-te | static]} | ospf | ospf-
extern1 | ospf-extern2 | ospf-inter | ospf-intra | rip | static |
ospf-inter-vr {unicast} {vr vrname}

```

Description

Displays the contents of the IP routing table for routes with the specified origin.

Syntax Description

auto-peering	Specifies <i>BGP</i> auto-peering routes.
bgp	Specifies <i>BGP</i> routes.

blackhole	Specifies blackhole routes.
bootp	Specifies BOOTP routes.
direct	Specifies direct routes.
ebgp	Specifies E-BGP routes.
embgp	Specifies EMBGP routes.
ibgp	Specifies I-BGP routes.
icmp	Specifies <i>ICMP</i> routes.
imbgp	Specifies IMBGP routes.
isis	Specifies IS-IS routes.
isis-level-1	Specifies IS-IS Level 1 routing.
isis-level-1-external	Specifies IS-IS Level 1 External routing.
isis-level-2	Specifies IS-IS Level 2 routing.
isis-level-2-external	Specifies IS-IS Level 2 External routing.
mbgp	Specifies MBGP routes.
mpls	Specifies <i>MPLS</i> routes. This option is available only on platforms that support the MPLS feature pack, which is described in the Switch Engine 32.3 Feature License Requirements document.
evpn	Specifies EVPN routes.
signaling-protocol [ldp rsvp-te static]	Specifies an MPLS signaling protocol. This option is available only on platforms that support the MPLS feature pack, which is described in the Switch Engine 32.3 Feature License Requirements document.
ospf	Specifies <i>OSPF</i> routes.
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 <i>RIP</i> routes.
static	Specifies static routes.
ospf-inter-vr	Specifies OSPF inter-VR routes.
unicast	Displays unicast routes with the specified origin.
vrname	Specifies a VR or VRF.

Default

N/A.

Usage Guidelines

None.

Example

The following example displays all the BGP routes:

```
# show iproute origin bgp
```

History

This command was first available in ExtremeXOS 10.1.

The **ipv4** keyword was added in ExtremeXOS 11.2.

The **embgp**, **mbgp**, and **mpls** options were added and the **ipv4** and **multicast** options removed in ExtremeXOS 12.2.2.

The **evpn** option was added in ExtremeXOS 30.7.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show iproute protection ping

```
show iproute {ipv4 | ipv6} protection ping {v4_or_v6_gateway} {vr
vr_name} {detail}
```

Description

Displays the current ping protection state for IPv4 or IPv6 routes, along with the number of protected routes. The information can be displayed per gateway, or for all gateways in a given VR.

Syntax Description

ipv4	Specifies IPv4 routes (default).
ipv6	Specifies IPv6 routes.
protection ping	Specifies protection ping information.
v4_or_v6_gateway	Name of the specific IPv4 or IPv6 gateway to show protected routes for.
vr <i>vr_name</i>	Shows all protected routes for the specified VR.
detail	Displays timestamps of the last time each gateway transitioned to the "up" state and to the "down" state.

Default

If not specified, IPv4 is the default.

Example

The following example shows all ping protected IPv4 routes only:

```
# show iproute protection ping
Ping protection for static IPv4 routes : Enabled
Interval : 2
Misses : 3
```

Gateway	State	# Routes	VR
1.2.3.4	Down	1	VR-Default
1.2.3.5	Up	1	VR-Default

The following example shows ping protected IPv6 routes only:

```
# show iproute ipv6 protection ping
Ping protection for static IPv6 routes : Enabled
Interval : 2
Misses : 3
```

Gateway	State	# Routes	VR
2000:2000:2000:2000:2000:2000:2001	Up	2	VR-Default
3000::2	Down	2	VR-Default

The following example shows ping protected IPv6 routes with gateway up/down transition timestamp information:

```
# show iproute ipv6 protection ping detail
Ping protection for static IPv6 routes : Enabled
Interval : 2
Misses : 3
```

```
Gateway: 2000:2000:2000:2000:2000:2000:2001
State: Up
# Routes: 2
VR: VR-Default
Last Up: Thu Mar 31 22:44:12 2016
Last Down: Thu Mar 31 22:42:22 2016
```

```
Gateway: 3000::2
State: Down
# Routes: 2
VR: VR-Default
Last Up: --
Last Down: Thu Mar 31 22:42:22 2016
```

History

This command was first available in ExtremeXOS 22.1.

Platform Availability

This command is available on all platforms with any license level as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

show iproute reserved-entries statistics

```
show iproute reserved-entries statistics {slot slot_num}
```

Description

Displays the current usage statistics of the Longest Prefix Match (LPM) hardware table and the Layer 3 hardware hash table by resource type.

Syntax Description

<code>slot_num</code>	For SummitStack only, this option displays the statistics for the specified slot.
-----------------------	---

Default

N/A.

Usage Guidelines

This command shows the current number of IP routes and local and remote IPv4 hosts in the LPM hardware table. It also shows the number of IPv4 unicast, multicast, and IPv6 unicast entries in the Layer 3 hardware hash table. The theoretical maximums for each individual resource type are shown at the bottom of the output. These maximum values cannot all be achieved simultaneously, and individual values might not be reached depending on the addresses or routes in use.

The ExtremeXOS software supports the coexistence of higher- and lower-capacity hardware. To allow for coexistence and increased hardware forwarding, when the number of IPv4 routes exceeds 25,000, the lower-capacity hardware automatically transitions from using LPM routing to forwarding of individual remote hosts, also known as IP Forwarding Database (IP *FDB*) mode. Higher-capacity hardware continues using LPM routing. Lower-capacity hardware operating in IP FDB mode will be indicated with a d flag in the output of `show iproute reserved-entries statistics` command, indicating that only direct routes are installed. For more information, see “Coexistence of Higher- and Lower-Capacity Hardware” in the [Switch Engine 32.3 User Guide](#).

Example

The following example displays usage statistics for the LPM and Layer 3 hardware tables:

```
# show iproute reserved-entries statistics
|-----In HW Route Table-----| |-----In HW L3
Hash Table-----|
IPv6   IPv4   IPv6
Slot  Type
Local  MCast  MCast
-----
1      5720-24MXW
0      0      0
# Used Routes  # IPv4 Hosts  IPv4  IPv4
IPv4  IPv6  Local Remote  Local  Rem.
-----
Theoretical maximum for each resource type:
5320 16-port, 5320 24-port      8160  4080   8192  8192  14334  16384  8192
* 8192 * 4096
5320 48-port, 5420F      12256  6128   12285  12288  14334  16384  8192
* 8192 * 4096
5420M      12256  6128   12288  12288  28670  32768  16384
*
```

```
16384 * 8192
5520                                16352  8176  16384  16384  41982  65536  18429
*
32768 * 16384
5720MW                                16352  8176  16384  16384  61438  98304  24573
*
49152 * 24576
5720MXW                                24544 12272  24576  24576 143358 163840  81920
*
81920 * 40960

Flags: (!) Indicates all reserved route entries in use.
        (d) Indicates only direct IPv4 routes are installed.
        (>) Some IPv6 routes with mask > 64 bits are installed and do not use
              entries in the internal HW Route Table.
        (*) Assumes IP Multicast compression is on.
        (M) IPMC entries stored in L2 MAC Table when lookup-key is 'mac-vlan'.
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show iproute reserved-entries

```
show iproute reserved-entries {slot slot_num}
```

Description

Displays the configured number of IPv4 and IPv6 routes reserved in the Longest Prefix Match (LPM) hardware table.

Syntax Description

slot_num	For SummitStack only, this option displays the reservations for the specified slot.
----------	---

Default

N/A.

Usage Guidelines

The IPv4 Routes column in the command output shows whether IPv4 routes are stored in internal or external LPM tables.

To modify the configuration that this command displays, use the following command:

```
configure iproute reserved-entries [ num_routes_needed | maximum |
default ] slot [all |slot_num]
```

Example

The following example displays the reserved space for IP routes:

```
# show iproute reserved-entries

Slot  Type                                     IPv4      # Reserved Routes      Minimum #
-----  -----                                     Routes    IPv4    (or IPv6)                IPv4 Hosts
1       5720-24MXW                                Internal  20432   ( 10216) [default]       4112

Maximum supported # IPv4 (or IPv6) Reserved Routes:
5320 16-port, 5320 24-port                Internal   8160   ( 4080)
5320 48-port, 5420F, 5420M                Internal  12256   ( 6128)
5520                                       Internal  16352   ( 8176)
5720MW                                    Internal  16352   ( 8176)
5720MXW                                   Internal  24544   (12272)

Note: IPv4 Hosts can occupy unused HW Route table space,
      unless # Reserved Routes is "maximum".
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show ip-security anomaly-protection notify cache ports

```
show ip-security anomaly-protection notify cache ports port_list
```

Description

Displays most anomaly notification caches.

Syntax Description

<code>port_list</code>	Specifies one or more ports or slots and ports.
------------------------	---

Default

N/A.

Usage Guidelines

This command displays most anomaly notification caches.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show ip-security arp gratuitous-protection

```
show ip-security arp gratuitous-protection
```

Description

If configured for gratuitous ARP, displays the gratuitous ARP protection configuration on the switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

The switch displays the name of each VLAN configured for gratuitous ARP.

If you do not have gratuitous ARP configured, the switch does not display any VLAN information.

Example

The following command displays the gratuitous ARP configuration on the switch:

```
show ip-security arp gratuitous-protection
```

The following is sample output from this command:

```
Gratuitous ARP Protection enabled on following VLANs:  
Default, test
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show ip-security arp learning

```
show ip-security arp learning [ {vlan} vlan_name | vlan vlan_list]
```

Description

Displays how the switch builds an ARP table and learns MAC addresses for devices on a specific VLAN and associated member ports.

Syntax Description

<i>vlan_name</i>	Specifies the name of the VLAN.
<i>vlan_list</i>	Specifies the name of the VLAN list of IDs.

Default

N/A.

Usage Guidelines

The switch displays the following ARP learning information:

- Port—The member port of the VLAN.
- Learn-from—The method the port uses to build the ARP table. The methods are:
 - ARP—ARP learning is enabled. The switch uses a series of requests and replies to build the ARP table.
 - DHCP—DHCP secured ARP is enabled. The switch uses DHCP snooping to build the ARP table.
 - None—Both DHCP secured ARP and ARP learning are disabled.

Example

The following sample output displays how the switch builds its ARP table for the VLAN learn:

```
# show ip-security arp learning vlan learn
Port          Learn-from
-----
2:1           ARP
2:2           DHCP,  poll 300 sec, retries 3
2:3           ARP
2:4           None
2:5           ARP
2:6           ARP
2:7           ARP
2:8           ARP
```

History

This command was first available in ExtremeXOS 11.6.

The *vlan_list* variable was added in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show ip-security arp validation

show ip-security arp validation [{**vlan**} *vlan_name* | **vlan** *vlan_list*]

Description

Displays ARP validation information for the specified VLAN.

Syntax Description

<i>vlan_name</i>	Specifies the name of the VLAN.
<i>vlan_list</i>	Specifies a VLAN list of IDs.

Default

N/A.

Usage Guidelines

The switch displays the following ARP validation information:

- Port—Indicates the port that received the ARP entry.
- Validation—Indicates how the entry is validated.
- Violation-action—Determines what action(s) the switch takes when an invalid ARP is received.

Example

The following sample output displays ARP validation on for the VLAN valid:

```
# show ip-security arp validation vlan valid
```

Port	Validation	Violation-action
7	DHCP	drop-packet, block-port for 120 seconds, snmp-trap
23	DHCP	drop-packet, block-port for 120 seconds, snmp-trap

The following example shows ARP validation if configured through RADIUS:

```
# show ip-security arp validation "Default"
ARP Validation enabled via RADIUS
```

Port	Validation	Violation-action
1	DHCP	drop-packet
2	DHCP	drop-packet
30	DHCP	drop-packet

History

This command was first available in ExtremeXOS 11.6.

The *vlan_list* variable was added in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show ip-security arp validation violations

```
show ip-security arp validation violations [ {vlan} vlan_name | vlan
vlan_list]ports [ports | all]
```

Description

Displays the violation count on an ARP validation.

Syntax Description

<i>vlan_name</i>	Specifies the name of the <u>VLAN</u> .
<i>vlan_list</i>	Specifies a VLAN list of IDs.
<i>ports</i>	Specifies the name of the port.
all	Specifies all ports.

Default

N/A.

Usage Guidelines

The switch displays the following ARP validation information:

- Port—Indicates the port that received the ARP entry.
- Validation—Indicates how the entry is validated.
- Violation count—Indicates the number of violations for each port.

Example

The following sample output displays ARP validation violation counts on all ports:

```
# show ip-security arp validation violations ragu ports all
-----
Port      Validation Violation Count
-----
1:1 ip,DHCP 1233
1:3 ip,DHCP 3425
1:4 ip,DHCP 5654
1:5 ip,DHCP 0
1:6 ip,DHCP 3645
```

History

This command was first available in ExtremeXOS 12.1.

The `vlan_list` variable was added in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show ip-security dhcp-snooping entries

```
show ip-security dhcp-snooping entries [ {vlan} vlan_name | vlan
vlan_list]
```

Description

Displays the *DHCP* bindings database on the switch.

Syntax Description

<i>vlan_name</i>	Specifies the name of the DHCP-snooping <i>VLAN</i> .
<i>vlan_list</i>	Specifies the DHCP-snooping VLAN list of IDs.

Default

N/A.

Usage Guidelines

The switch displays the following DHCP bindings database information:

- VLAN—The name of the DHCP-snooping VLAN.
- IP Addr—The IP address of the untrusted interface or client.
- MAC Addr—The MAC address of the untrusted interface or client.
- Port—The port number where the untrusted interface or client attempted to access the network.

Example

The following sample output displays the DHCP bindings database on the switch:

```
# show ip-security dhcp-snooping entries vlan dhcpVlan
-----
Vlan: dhcpVlan
-----
Server      Client
IP Addr      MAC Addr      Port      Port
-----
172.16.100.9  00:90:27:c6:b7:65  1:1      1:2
```

History

This command was first available in ExtremeXOS 11.6.

The `vlan_list` variable was added in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show ip-security dhcp-snooping information circuit-id port-information

```
show ip-security dhcp-snooping information circuit-id port-information
    ports [port_list | all ]
```

Description

Displays the port information portion of the circuit ID for the indicated port(s).

Syntax Description

<code>port_list</code>	Specifies one or more ports.
all	Specifies all ports.

Default

N/A.

Usage Guidelines

This command displays the port information portion of the circuit ID for the indicated ports.

Example

The following command:

```
# show ip-security dhcp-snooping information circuit-id port-information ports 1-7
```

Displays the following output:

```
Port          Circuit-ID Port information string
----          -
1             portinfostring1
2             portinfostring2
3             portinfostring3
4             portinfostring4
5             portinfostring5
Port          Circuit-ID Port information string
----          -
6             1006
7             1007
Note: The full Circuit ID string has the form '<Vlan Info>--<Port Info>'
```

History

This command was first available in ExtremeXOS 12.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show ip-security dhcp-snooping information-option

```
show ip-security dhcp-snooping information-option
```

Description

Displays the *DHCP* relay agent option (option 82) settings.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command displays DHCP relay agent option (option 82) settings. For example, the following command:

```
show ip-security dhcp-snooping information-option
```

Generates the following output:

```
Information option insertion: Enabled Information option checking : Disabled Information  
option policy : Drop
```

The following command:

```
show ip-security dhcp-snooping information-option
```

Generates the following output:

```
Information option insertion: Enabled Information option checking : Enabled Information  
option policy : Keep
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show ip-security dhcp-snooping information-option circuit-id vlan-information

```
show ip-security dhcp-snooping information-option circuit-id vlan-  
information {{vlan} [vlan_name | vlan_list]}
```

Description

Displays the VLAN information portion of the circuit ID for the indicated VLAN.

Syntax Description

<i>vlan_name</i>	Specifies a VLAN name.
<i>vlan_list</i>	Specifies a VLAN list of IDs.

Default

N/A.

Usage Guidelines

This command displays the VLAN information portion of the circuit ID for the indicated VLAN. When a VLAN is not specified, the circuit ID information for all the VLANs is displayed.

Example

The following sample output displays circuit ID information for vlan "Mktg":

```
# show ip-security dhcp-snooping information-option circuit-id vlan-information vlan Mktg  
Vlan                               Circuit-ID vlan information string  
----                               -  
Mktg                               DSLAM1  
Note: The full Circuit ID string has the form <Vlan Info>-<Port ifIndex>.
```

History

This command was first available in ExtremeXOS 12.1.

The *vlan_list* variable was added in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show ip-security dhcp-snooping information remote-id

```
show ip-security dhcp-snooping information remote-id
```

Description

Shows the *DHCP* relay agent remote ID.

Syntax Description

This command has no arguments or variables.

Default

If a remote ID has not been configured, or it has been unconfigured, the default remote ID is the switch's MAC address. However, this default remote ID does not appear in this command.

Usage Guidelines

This command shows the remote ID set by the `configure ip-security dhcp-snooping information remote-id [system-name | remote-id_info]` command. If the remote ID has never been configured, or it has been unconfigured (`unconfigure ip-security dhcp-snooping information remote-id`), the default remote ID is the switch's MAC address, which does not appear in the output of this command.

Example

The following command shows the DHCP remote ID:

```
# show ip-security dhcp-snooping information-option remote-id

Switch          Remote-ID information string
-----
X465-48P        mydhcp
```

The following command shows that the DHCP remote ID has never been configured, or that it was unconfigured, and that the default (switch MAC address) is being used, but this default remote ID does not appear in the output of the command:

```
Switch          Remote-ID information string
-----
Note: If neither System Name nor the Customized remote id is configured
      the default will be the Switch MAC address, that will not be shown
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show ip-security dhcp-snooping

```
show ip-security dhcp-snooping [ {vlan} vlan_name | vlan vlan_list]
```

Description

Displays the *DHCP* snooping configurations on the switch.

Syntax Description

<i>vlan_name</i>	Specifies the name of the DHCP-snooping <i>VLAN</i> .
<i>vlan_name</i>	Specifies the name of the DHCP-snooping VLAN list.

Default

N/A.

Usage Guidelines

The switch displays the following DHCP snooping information:

- DHCP snooping enabled on ports—The ports that have DHCP snooping enabled.
- Trusted ports—The ports configured as trusted ports.
- Trusted DHCP servers—The servers configured as trusted DHCP servers.
- Port—The specific port that has DHCP snooping enabled.
- Violation-action—The action the switch takes upon detecting a rogue DHCP packet on the port.

Example

The following sample output displays the DHCP snooping settings for the switch:

```
# show ip-security dhcp-snooping vlan "Default"
DHCP Snooping enabled on ports: 7, 9, 11
Trusted Ports: None
Trusted DHCP Servers: None
Bindings Restoration      : Enabled
Bindings Filename        : dhcpsonia.xsf
Bindings File Location   :
Primary Server   : 10.1.1.14, VR-Default, TFTP
Secondary Server: None
Bindings Write Interval : 5 minutes
Bindings last uploaded at:
-----
Port          Violation-action
-----
7             none
9             none
11            none
```

The following sample output displays the DHCP snooping settings for the switch if configured through RADIUS:

```
# show ip-security dhcp-snooping "Default"
DHCP Snooping enabled via RADIUS
DHCP Snooping enabled on ports: 1, 2, 30
Trusted Ports: None
Trusted DHCP Servers: 1.1.1.1, 1.1.2.1
Bindings Restoration      : Disabled
Bindings Filename        :
Bindings File Location    :
        Primary Server   : None
        Secondary Server : None
Bindings Write Interval   : 30 minutes
-----
Port          Violation-action
-----
1             drop-packet
2             drop-packet
30            drop-packet
```

History

This command was first available in ExtremeXOS 11.6.

The `vlan_list` variable was added in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show ip-security dhcp-snooping violations

```
show ip-security dhcp-snooping violations [ {vlan} vlan_name | vlan
vlan_list]
```

Description

Displays the MAC addressed from which the rouge DHCP packet was received by the switch.

Syntax Description

<i>vlan_name</i>	Specifies the name of the DHCP-snooping <u>VLAN</u> .
<i>vlan_list</i>	Specifies the DHCP-snooping VLAN list of IDs.

Default

N/A.

Usage Guidelines

The switch displays the following DHCP snooping information:

- Port—The specific port that received the rouge DHCP packet.
- Violating MAC—The MAC address from which the rouge DHCP was received by the switch.

Example

The following sample output displays the DHCP snooping violations for the VLAN green:

```
# show ip-security dhcp-snooping violations green
Violations seen on following ports
-----
Port                Violating MAC
-----
2:3                 00-0c-11-a0-3e-12
```

History

This command was first available in ExtremeXOS 11.6.

The `vlan_list` variable was added in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show ip-security source-ip-lockdown

```
show ip-security source-ip-lockdown
```

Description

Displays the source IP lockdown configuration on the switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

The switch displays the following source IP lockdown information:

- Port—Indicates the port that has DHCP snooping enabled and is configured for source IP lockdown.
- Locked IP Address—Indicates a valid DHCP-assigned address obtained by a DHCP snooping-enabled port or an authenticated static IP address.

Example

The following command displays the source IP configuration on the switch:

```
show ip-security source-ip-lockdown
```

The following is sample output from this command:

```
Ports          Locked IP Address
23 10.0.0.101
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show ipstats ipv6

```
show ipstats ipv6 [{vlan name | tunnel tunnel_name}]
```

Description

Displays IPv6 statistics handled by the CPU for the switch for a particular VLAN.

Syntax Description

<i>name</i>	Specifies a VLAN name.
<i>tunnel_name</i>	Specifies a tunnel name.

Default

N/A.

Usage Guidelines

This command only shows statistics of the CPU-handled packets. Not all packets are handled by the CPU. For example, packets forwarded in hardware do not increment the statistics counters.

Example

The following example displays IPv6 statistics for the VLAN accounting:

```
show ipstats ipv6 vlan accounting
```

History

This command was first available in ExtremeXOS 11.2.

The VR option was removed in ExtremeXOS 30.1.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

show ipstats

```
show ipstats {ipv4} [{vlan name } | {tunnel} tunnel_name]
```

Description

Displays IP statistics for the switch CPU for a particular VLAN.

Syntax Description

<i>name</i>	Specifies a VLAN name.
<i>tunnel_name</i>	Specifies a tunnel 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 the following tables.

Table 25: 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.

Table 25: Router Interface Statistics Field Definitions (continued)

Field	Definition
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 example displays IP statistics for the VLAN "accounting":

```
show ipstats vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

The **ipv4** keyword was added in ExtremeXOS 11.2.

The VR option was removed in ExtremeXOS 30.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show ipv6 dad

```
show ipv6 dad [{vr} vr_name {ip_address} | vr all | [ { vlan }
  vlan_name | vlan vlan_list ] {tentative | valid | duplicate} | [{vr}
  vr_name} ipaddress] | {tunnel} tunnel_name
```

Description

Displays the configuration and run time status for the DAD feature on the specified IPv6 interface.

Syntax Description

<i>vr_name</i>	Specifies a VR for which to display the DAD information.
<i>ip_address</i>	Specifies an IPv6 address for which to display the DAD information.
<i>vlan_name</i>	Specifies a VLAN name for which to display the DAD information.
<i>vlan_list</i>	Specifies a VLAN list of IDs to display the DAD information.
tentative	Displays information for IPv6 interfaces for which the status is up and the DAD check is incomplete.

valid	Displays information for IPv6 interfaces for which the status is up, the DAD check is complete, and no duplicate IPv6 addresses were detected.
duplicate	Displays information for IPv6 interfaces for which the status is down because a duplicate IPv6 address was detected.

Default

If you do not specify a VR or VRF, the command applies to the current VR context.

Usage Guidelines

The **vr all** option displays DAD information for all IPv6 interfaces on the switch.

Example

The following example displays the DAD feature status for all interfaces in the current VR context:

```
IPv6 Duplicate Address Detection

DAD Status           : On
Max Solicitation Attempts : 1

Virtual Router      Interface      Flags Failures IP
Address
-----
VR-Default          vwest          D1E-U          1
fe80::204:96ff:fe99:1424%vwest
VR-Default          vwest          -1E-U          0
fe80::1%vwest
VR-Default          vwest          --E-U          0 2010::1

Flags: (D) Duplicate address detected, (E) L2 Interface enabled,
       (l) Link-local address, (L) Loopback enabled,
       (t) Tentative due to the lack of a valid link-local address,
       (T) Tentative address, (U) L2 Interface up
```

History

This command was first available in ExtremeXOS 12.6.

The `vlan_list` variable was added in ExtremeXOS 16.1.

The Conflict MAC heading was removed and Interface Failures heading added; the P (prefix address) flag was removed and the l (Link-local address) was added; the t flag was renamed (tentative due to the lack of a valid link-local address); and the **detail** option was removed in ExtremeXOS 30.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show iqagent

show **iqagent discovery**

Description

Displays information about the ExtremeCloud™ IQ Agent.

Syntax Description

iqagent	Specifies showing information about IQ Agent.
discovery	Specifies showing information about IQ Agent discovery.

Default

N/A.

Usage Guidelines

If IQ Agent is not running, only the first section appears. A device is registered on an IQ server using the serial number.

Example

The following example shows information about IQ Agent:

```
# show iqagent
-----
State                Enabled
Active VR            VR-Mgmt
Detected VR          VR-Mgmt
XIQ Address           10.16.231.95
Serial Number         1405G-00056
-----
```

The following example shows discovery information for IQ Agent:

```
# show iqagent discovery
-----
Virtual-Router  : VR-Default
  XIQ FQDN CLI   :
  XIQ DHCP IP    :
  XIQ DHCP FQDN : test.com
  Redirector     :
-----
Virtual-Router  : VR-Mgmt
  XIQ FQDN CLI   : 10.16.231.95
  XIQ DHCP IP    : 1.2.3.4
  XIQ DHCP FQDN : test2.com
  Redirector     :
-----
```

The following example shows information about IQ Agent with a configured HTTP proxy (line 14):

```
# show iqagent
-----
State                               Enabled
Active VR                          VR-Mgmt
Detected VR                         VR-Mgmt
XIQ Address                         10.51.3.175
Serial Number                       1546N-40030
-----
Status                             CONNECTED TO XIQ
Version                             0.4.5
Last Onboard Time                   15:05:38 11 12 2020 UTC
Debug                               ON
Source Interface                     DEFAULT
HTTP Proxy                          iqagent@10.51.3.163:3128
-----
Association Method                   CLI
Association Url                      https://10.51.3.175/hac-webapp/rest/v1/association
Poll URL                            https://10.51.3.175/hac-webapp/rest/v1/poll/
1546N-40030
-----
Last Health Status                  SUCCESS
Last Health Success Time             15:05:41 11 12 2020 UTC
-----
Last Poll Status                    SUCCESS
Last Poll Success Time               15:05:38 11 12 2020 UTC
-----
Last Monitor Status                  SUCCESS
Last Monitor Success Time             15:05:48 11 12 2020 UTC
-----
```

The following example shows information when there is a configured IQ Agent Management VLAN:

```
# show iqagent
-----
State                               Enabled
Active VLAN                         Mgmt
User VR                             VR-Mgmt
User VLAN                           Mgmt
XIQ Address                         10.51.3.175
Serial Number                       1409G-00236
-----
Status                             CONNECTED TO XIQ
Version                             922.755.1
Last Onboard Time                   09:57:48 05 10 2022 UTC
Debug                               ON
Source Interface                     10.127.2.117
HTTP Proxy                          None
-----
Association Method                   CLI
Association Url                      https://10.51.3.175/hac-webapp/rest/v1/association
Poll URL                            https://10.51.3.175/hac-webapp/rest/v1/poll/
1409G-00236
-----
Last Health Status                  SUCCESS
Last Health Success Time             10:13:50 05 10 2022 UTC
-----
Last Poll Status                    SUCCESS
Last Poll Success Time               10:13:48 05 10 2022 UTC
-----
Last Monitor Status                  SUCCESS
Last Monitor Success Time             10:13:49 05 10 2022 UTC
-----
```

History

This command was first available in ExtremeXOS 30.7.

IQ Agent status was added in ExtremeXOS 31.1.

HTTP proxy support was added in ExtremeXOS 31.3.

IQ Agent Management VLAN was added in Release 32.1.

Platform Availability

This command is available on ExtremeSwitching 5420 and 5520 series switches.

show iqagent discovery detail

```
show iqagent discovery detail {vr vr_name}
```

Description

Displays detailed information about a Virtual Router connected to the ExtremeCloud™ IQ Agent.

Syntax Description

iqagent	Specifies showing information about IQ Agent.
discovery	Specifies showing information about IQ Agent discovery.
detail	Specifies showing a detailed view of connection to ExtremeCloud IQ.
vr	Specifies showing a detailed view for a specific Virtual Router.
<i>vr_name</i>	Specifies the Virtual Router name.

Default

N/A.

Usage Guidelines

If IQ Agent is not running, only the first section displays. A device is registered on an IQ server using the serial number.

Example

The following example shows information about a configured static IQ Agent (10.51.3.171) and DHCP providing both sub-options 225 and 226:

```
# show iqagent discovery detail
IQ Agent configured on VR VR-Default
  HTTP proxy                : Not configured
  User configured FQDN/IP    : 10.51.3.171
  XIQ Server Status          : Connected
  DHCP opt 43, sub-opt 225 (FQDN) : fqdn.xiq-ut.com
```



```

Resolved IP           : 10.51.3.175
XIQ Server Status    : Connected
DHCP opt 43, sub-opt 226 (IP) : 10.16.231.94
XIQ Server Status    : Connected
Redirector           : hac.extremecloudiq.com
Resolved IP           : Unresolved

```

The following example shows information about a configured static IQ Agent (10.51.3.171) onboarded device with hac.extremecloudiq.com that is reachable:

```

# show iq discovery detail
IQ Agent configured on VR VR-Default
  HTTP proxy           : Not configured
  User configured FQDN/IP : 10.51.3.171
  XIQ Server Status    : Connected
  DHCP opt 43, sub-opt 225 (FQDN) : Not provided
  DHCP opt 43, sub-opt 226 (IP) : Not provided
  Redirector           : hac.extremecloudiq.com
  Resolved IP          : 34.253.190.222
  Cloud IQ Status      : Connected
  Data Center [1848F-10070] : va2.extremecloudiq.com
  Resolved IP          : 34.202.197.10
  XIQ Server Status    : Connected

```

The following example shows a user configured FQDN that is not resolvable:

```

show iq discovery detail
IQ Agent configured on VR VR-Default
  HTTP proxy           : Not configured
  User configured FQDN/IP : xiq.xiq-ut.com
  Resolved IP          : Unresolved
  DHCP opt 43, sub-opt 225 (FQDN) : Not provided
  DHCP opt 43, sub-opt 226 (IP) : 10.16.231.94
  XIQ Server Status    : Connected
  Redirector           : hac.extremecloudiq.com
  Resolved IP          : 34.253.190.222
  Cloud IQ Status      : Connected
  Data Center [1848F-10070] : va2.extremecloudiq.com
  Resolved IP          : 34.202.197.10
  XIQ Server Status    : Connected

```

History

This command was first available in ExtremeXOS 31.4.

Platform Availability

This command is available on ExtremeSwitching 5420 and 5520 series switches.

show isis

```
show isis
```

Description

This command displays the global IS-IS configuration information as well as a summarized router process listing.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

The displayed global configuration information includes the restart enablement, restart grace period, and import-policy setting. The router process listing includes the area name, system ID, whether it's enabled, the IS type, and a count of associated interfaces and area addresses. This command applies only to the IS-IS router processes running in the current virtual router.

Example

The following command displays IS-IS information:

```
show isis
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Premier license as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

show isis area summary-addresses

```
show isis area area_name summary-addresses
```

Description

This command displays the configured IPv4 and IPv6 summary addresses for the specified area.

Syntax Description

<i>area_name</i>	Specifies the router process for which you want to display summary addresses.
------------------	---

Default

N/A.

Usage Guidelines

None.

Example

The following command displays the summary addresses for areax:

```
show isis area areax summary-addresses
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Premier license as described in the [**Switch Engine 32.3 Feature License Requirements**](#) document.

show isis area

```
show isis area [area_name | all]
```

Description

This command displays configuration information for a specific router process or for all IS-IS router processes.

Syntax Description

<i>area_name</i>	Specifies the router process for which to display IS-IS area information.
all	Displays information for all IS-IS router processes.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays IS-IS configuration for areax:

```
show isis area areax
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Premier license as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

show isis counters

```
show isis counters {area [area_name | all] | vlan [vlan_name | all]}
```

Description

This command displays counters for an area or a [VLAN](#).

Syntax Description

<i>area_name</i>	Specifies a router process for which to display the IS-IS counters. If you do not specify an IS-IS area, the software displays counters for all areas.
<i>vlan_name</i>	Specifies a VLAN for which to display IS-IS counters.
all	Displays IS-IS counters for all areas or VLANs.

Default

None.

Usage Guidelines

If you enter the show isis counters command without any additional keywords or parameters, the software displays the counters for all areas.

Example

The following command displays the IS-IS counters for the configured area:

```
show isis counters
```

The following command displays the IS-IS counters for the SJvlan VLAN:

```
show isis counters vlan SJvlan
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Premier license as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

show isis lsdb

```
show isis lsdb {area area_name {lsp-id lsp_id}} {level [1|2]} {detail | stats}
```

Description

Displays a summary of the IS-IS link state database for one or all IS-IS router processes running in the current virtual router.

Syntax Description

<i>area_name</i>	Specifies the name of a router process for which to display the IS-IS link state database. If the area name is omitted, this command displays information for all areas.
<i>lsp_id</i>	Limits the display to the specified LSP ID, which is specified in the form <i>system id.pseudonode ID- LSP number</i> . For example: 0102.0ff2.0023.00-01. The pseudonode ID and LSP numbers are optional. If they are not included, multiple listings might appear.
level [1 2]	Limits the display to LSPs for either level 1 or level 2.
detail	Expands the display to include the LSP TLVs and IPv4 and IPv6 reachability information. The displayed information varies depending on what is included in the LSPs.
stats	Displays counts of the number of LSPs and prefixes in the LSP database.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays information for a specific LSP:

```
show isis lsdb area areax lsp-id 0102.0ff2.0023.00-01
```

The following example shows the display for the stats option:

```
(debug) Switch.6 # show isis lsdb stats
```

```
Area "a1" :
IS-IS Level-1 Link State Database:
LSPs (including fragments)      : 4
Internal Prefixes (Type 128)    : 7
External Prefixes (Type 130)    : 0
IPv4 Prefixes (Type 135)       : 0
IPv6 Prefixes (Type 236)       : 0
MT IPv4 Prefixes (Type 235)    : 0
MT IPv6 Prefixes (Type 237)    : 0
IS-IS Level-2 Link State Database:
LSPs (including fragments)      : 1
Internal Prefixes (Type 128)    : 5
External Prefixes (Type 130)    : 0
IPv4 Prefixes (Type 135)       : 0
IPv6 Prefixes (Type 236)       : 0
MT IPv4 Prefixes (Type 235)    : 0
MT IPv6 Prefixes (Type 237)    : 0
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Premier license as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

show isis neighbors

```
show isis neighbors {area area_name} {vlan vlan_name} {ipv4 | ipv6}
                    {detail}
```

Description

This command displays information about neighbors and their adjacencies.

Syntax Description

area_name	Specifies a router process for which to specify neighbor information. If you do not specify an IS-IS area, the software displays the neighbors in all areas.
vlan_name	Specifies a VLAN for which to specify neighbor information.
ipv4	Displays only the neighbors that advertise the IPv4 protocol as supported.
ipv6	Displays only the neighbors that advertise the IPv6 protocol as supported.
detail	Displays detailed information for IS-IS neighbors.

Default

N/A.

Usage Guidelines

If you do not specify either the `ipv4` or the `ipv6` keyword, this command displays all neighbors regardless of the supported protocol.

Example

The following command displays IS-IS neighbor information for areax:

```
show isis neighbors area areax
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Premier license as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

show isis topology

```
show isis topology {area area_name {level [1 | 2]}} {ipv4 | ipv6}
```

Description

This command displays the topology for IPv4, IPv6, or both IPv4 and IPv6 for the specified area and level.

Syntax Description

<i>area_name</i>	Specifies the router process for which the topology applies.
level [1 2]	Specifies the IS-IS level of the topology you want to view.
ipv4	Selects the IPv4 topology for display. If you omit the <code>ipv4</code> and <code>ipv6</code> options, the IPv4 topology appears.
ipv6	Selects the IPv6 topology for display. If you omit this option, the IPv4 topology appears.

Default

None.

Usage Guidelines

Each known IS in the area or domain is displayed along with the next-hop and metric information.

Example

The following command display IPv4 topology information for areax:

```
show isis topology area areax ipv4
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Premier license as described in the [***Switch Engine 32.3 Feature License Requirements***](#) document.

show isis vlan

```
show isis vlan {enabled | { vlan_name | all } }
```

Description

This command displays configuration and status information about the specified IS-IS interface.

Syntax Description

enabled	Displays information only for IS-IS-enabled <i>VLANs</i> .
<i>vlan_name</i>	Specifies a VLAN for which to display IS-IS interface information.
all	Displays information on all IS-IS interfaces.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays IS-IS interface information for the SJvlan VLAN:

```
show isis vlan SJvlan
ISIS Interfaces Summary :
-----
VLAN      Area      State    Cfg      Address
-----
v2         a1         u46--    p1246     2.1.1.2/24
2001:db8:2010::1/64
v3         a1         u46-g    p1246     2001:db8:2011::2/64
```



```

State Flags :
u - Links up, d - Links down,
4 - IPv4 forwarding enabled, 6 - IPv6 forwarding enabled,
n - Multinetted (v4), g - Multiple global addresses (v6)
Cfg Flags   :
b - Broadcast interface, p - Point-To-Point interface,
1 - L1 circuit type, 2 - L2 circuit type,
4 - ISIS-enabled for IPv4, 6 - ISIS-enabled for IPv6

```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Premier license as described in the [***Switch Engine 32.3 Feature License Requirements***](#) document.

show keychain

```
show keychain {keychain_name {detail}}
```

Description

Displays the configured keychains.

Syntax Description

<i>keychain_name</i>	Defines a name for the new keychain. The range is 1-31.
detail	Specifies a detailed output.

Default

N/A.

Usage Guidelines

Use this command to display the configured keychains.

Example

The following is an example of the command's output with no specified keychain name:

```

# show keychain

```

Name	# Keys	# Clients	Accept Tolerance
auth1	3	2	10s
auth2	4	0	0s
ospfv3-keys	6	1	0s

Example

The following is an example of the command's output with a keychain name of keys1:

```
show keychain ospfv3-keys1
Accept Tolerance : 0s
Number of Clients: 1
```

Key ID	Algorithm	Start Time	End Time	Duration
+1	HMAC-SHA-256	2021-06-01 00:00:00	2021-07-01 00:00:00	2592000s
2	HMAC-SHA-512	2021-07-01 00:00:00	2021-08-01 00:00:00	2678400s
3	HMAC-SHA-256	2021-08-01 00:00:00	2021-09-01 00:00:00	2678400s

Legend: (*) Active Key, (+) Selected Key, (Z) Time is in UTC.

Example

The following is an example of the command's output with a keychain name of keys1 with detail specified:

```
show keychain ospfv3-keys1 detail
Accept Tolerance : 0s
Number of Clients : 1
Clients : ospfv3-2
```

Key ID : +1
 Encrypted Key String: #y209VBgcxFr4NkVKnadBL5y1s9cvnA==
 Algorithm : HMAC-SHA-256
 Start Time : 2021-06-01 00:00:00
 End Time : 2021-07-01 00:00:00
 Duration : 2592000s

Key ID : 2
 Encrypted Key String: #f4Vtob4C5TdAs9Qf0w/pbB8RecE9vA==
 Algorithm : HMAC-SHA-512
 Start Time : 2021-07-01 00:00:00
 End Time : 2021-08-01 00:00:00
 Duration : 2678400s

Key ID : 3
 Encrypted Key String: #bvE/wSf/6sztYL4wJw8ZUTWjJT2ZWw==
 Algorithm : HMAC-SHA-256
 Start Time : 2021-08-01 00:00:00
 End Time : 2021-09-01 00:00:00
 Duration : 2678400s

Legend : (*) Active Key, (+) Selected Key.

History

This command was first available in ExtremeSwitching 31.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show l2pt

```
show l2pt
```

Description

Displays the global parameters for L2PT.

Syntax Description

l2pt	Displays global Layer 2 protocol tunneling parameters.
-------------	--

Default

Disabled.

Usage Guidelines

Use this command to display the global parameters for L2PT.

Example

The following is an example of the command's output:

```
# show l2pt
Encapsulation Destination MAC Address: 01:00:00:01:01:02
```

History

This command was first available in ExtremeXOS 15.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show l2pt profile

```
show l2pt profile profile_name
```

Description

Displays the contents of an L2PT profile.

Syntax Description

profile	Displays profile that defines L2PT configuration for L2 protocols.
<i>profile_name</i>	Displays only the specified profile.

Default

Disabled.

Usage Guidelines

Use this command to display the contents of an L2PT profile.

Example

The following is an example of the command's output:

```
# show l2pt profile

Profile Name          Protocol Filter Name          DSCP
Replace              -----
-----
my_l2pt_access_prof   my_list                       tunnel  1      63
Yes                  my_other_list                 tunnel  7      0
No
my_l2pt_network_prof  mylist                        encap   20
No                  my_other_list                 encap   8
No                  my_none~_list                 none    25
No

# show l2pt profile my_l2pt_access_prof

DSCP
Profile Name          Protocol Filter Name          Action  CoS    DSCP
Replace              -----
-----
my_l2pt_access_prof   my_list                       tunnel  1      23
Yes                  my_other_list                 tunnel  7      0
No
```

History

This command was first available in ExtremeXOS 15.5.

Support for DSCP was added in ExtremeXOS 31.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show L2stats

```
show L2stats [ { vlan } vlan_name | vlan vlan_list ]
```

Description

Displays the counters for the number of packets bridged, switched, and snooped (Layer 2 statistics).

Syntax Description

<i>vlan_name</i>	Specifies a <u>VLAN</u> name.
<i>vlan_list</i>	Specifies a VLAN list of IDs.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays the counters for the number of packets bridged, switched, and snooped (Layer 2 statistics) for the VLAN accounting:

```
show L2stats accounting
```



Note

You can also enter the command as `show l2stats`. We use the uppercase letter here to avoid confusion with the numeral 1.

History

This command was first available in ExtremeXOS 11.0.

The *vlan_list* variable was added in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show l2vpn

```
show [ {l2vpn} vpls {vppls_name} | l2vpn vpws {vpws_name} | l2vpn ] {peer  
    ipaddress} {detail} | summary }
```

Description

Displays Layer 2 VPN configuration and status information.

Syntax Description

l2vpn	Displays specified Layer 2 VPN information.
<i>vpls_name</i>	Displays information for the specified VPLS.
<i>vpws_name</i>	Displays information for the specified VPWS.
<i>ipaddress</i>	Specifies a Layer 2 VPN peer for which to display information.
detail	Displays additional information in comprehensive detail format.
summary	Displays summary information about all VPLS or VPWS instances.

Default

N/A.

Usage Guidelines

The `show l2vpn` command (without any optional parameters) displays all currently configured Layer 2 VPN instances for the switch. The summarized list of Layer 2 VPN instances is displayed in alphabetical order based on the Layer 2 VPN name. Peers are displayed in the reverse of the order they were added.

When you specify a Layer 2 VPN peer, the display includes a list of all PWs established to the peer, the PW status and PW ID, and information about each Layer 2 VPN to which this peer belongs.

The following table describes the display fields that appear when this command is entered with the **detail** option.

Table 26: Selected show l2vpn Field Definitions

Field	Definition
VPLS Name	VPLS instance or domain name.
VPN ID	Virtual Private Network identifier.
Source Address	Source IP address.
VCCV Status	Virtual Circuit Connectivity Verification (VCCV) feature status, which is either Enabled or Disabled.
VCCV Interval Time	Displays the configured VCCV interval time.
VCCV Fault Multiplier	Displays the configured VCCV fault multiplier.
Redundancy Type	Displays the configured VPLS redundancy type, which is <i>EAPS</i> , <i>ESRP</i> , <i>STP</i> , or None.
Service Interface	Displays a <i>VLAN</i> or VMAN interface name.
Admin State	Displays the administrative state of the VPLS, which is either Enabled or Disabled.
Oper State	Displays the operational state of the VPLS, which is either Enabled or Disabled.
MTU	Displays the maximum transmission unit (MTU) size for the VPLS.

Table 26: Selected show l2vpn Field Definitions (continued)

Field	Definition
Ethertype	Displays the ethernet type for the service interface.
.1q tag	Displays the 802.1q priority tag for the VPLS.
Peer IP	Displays the IP address for the VPLS peer.
PW State	<p>PW State represents the state, or status, of a PW. The possible PW state values are:</p> <ul style="list-style-type: none"> • UP—The PW is fully operational and installed in hardware. Traffic is forwarded over PW and VPLS service VLAN/VMAN. • Down—The PW is not operational and is not installed in hardware. This only happens when the VPLS instance is disabled, VPLS service is disabled, or there is no service VLAN assigned to the VPLS. No traffic is forwarded. • Sgnl—The PW is in a signalling state. The PW is not operational, and no traffic is forwarded. This can occur for a number of reasons, including: No LDP adjacency to peer, No transport LSP to peer, No VC LSP to peer. • Remote peer not configuredReady—The PW has been signalled, but it has not been installed in hardware. Traffic is not forwarded. The PW can be in a Ready state for a number of reasons, including: <ul style="list-style-type: none"> ◦ The VPLS instance is configured for EAPS redundancy, and the EAPS shared port associated with this VPLS instance is Connected. ◦ The VPLS instance is configured for ESRP redundancy, and the ESRP domain associated with this VPLS instance is Slave. ◦ The service VLAN associated with this VPLS instance is down. ◦ The remote peer has signalled that it has a fault (remote PW status). The remote peer may have a fault due to its service VLAN being down.
PW Uptime	PW Uptime is the elapsed time that the PW has been in the UP state.
PW Installed	PW Installed is a flag to indicate whether the PW is installed in hardware or not. If the PW is in the UP state, this field is True, otherwise, this field is False.
Local PW Status	<p>Local PW Status displays the VC status of the local PW. The values are:</p> <ul style="list-style-type: none"> • No Faults—No faults detected. • PW-Tx—Local PSN-facing PW transmit fault. This is set if there is a problem with the VPLS transport LSP. • PW-Rx—Local PSN-facing PW receive fault. This is set if there is a problem with the VPLS transport LSP. • Att-Tx—Local attachment circuit transmit fault. This is set if there is a problem with the VPLS service VLAN. • Att-Rx—Local attachment circuit receive fault. This is set if there is a problem with the VPLS service VLAN. • Not Forwarding—The local PW is not forwarding. Look for more information in the PW State field. For example, if VPLS is configured for EAPS redundancy, the Local PW Status is Not Forwarding and the PW State is Ready whenever the EAPS Shared Port state is Connected.
Remote PW Status	Remote PW Status is the VC status of the remote PW. The values for this field are the same values as for Local PW Status.

Table 26: Selected show l2vpn Field Definitions (continued)

Field	Definition
PW Mode	PW Mode describes how the PW was configured. The values are: <ul style="list-style-type: none"> Core-to-Core—This VPLS instance is a core node, and the other end of the PW connects to a core node. Core-to-Spoke—This VPLS instance is a core node, and the other end of the PW connects to a spoke node. This is for HVPLS. Spoke-to-Core—This VPLS instance is a spoke node, and the other end of the PW connects to a core node. This is for HVPLS.
Transport LSP	Transport LSP is the LSP that is used to forward frames over the PW. When an LDP LSP is used as a transport, the display shows LDP LSP (Not configured). If an RSVP LSP is used, the name of the RSVP LSP being used as a transport LSP is displayed. An RSVP LSP can be specified as the LSP to use during VPLS configuration.
Next Hop I/F	Displays the interface name for the next hop router.
Next Hop Addr	Displays the interface IP address for the next hop router.
PW Rx Label	Receive label for the VPLS PW.
PW Rx Pkts	Total packets received on the VPLS PW.
PW Rx Bytes	Total bytes received on the VPLS PW.
Tx Label	Transmit label for the LSP.
PW Tx Label	Transmit label for the VPLS PW.
PW Tx Pkts	Total packets transmitted on the VPLS PW.
PW Tx Bytes	Total bytes transmitted on the VPLS PW.

The **l2vpn** keyword is introduced in ExtremeXOS Release 12.4 and is required when displaying VPWS information. For backward compatibility, the **l2vpn** keyword is optional when displaying VPLS information. However, this keyword will be required in a future release, so Extreme Networks recommends that you use this keyword for new configurations and scripts.

Example

The following example shows the display that appears when you enter the `show l2vpn` command without any options:

```
# show l2vpn
```

L2VPN Name	VPN ID	Flags	Services Name	Peer IP	State	Flags
Pws-3344	20	EAX--W	NONE			
jwcvpls	99	EAX--L	torix	11.100.100.219	Up	C--NV-
keeper	90	EAX--L	NONE			
pws-1	2009	EAX--W	pwserve	11.100.100.219	Up	----V-
pws-10	70	EAX--W	NONE			
pws-2	2008	EAX--W	pw2serve	11.100.100.219	Up	---NV-
pws-3	2007	EAX--W	NONE			
sarsparilla	80	EAX--W	NONE			
whoopwoo	100	EAX--L	NONE	11.100.100.219	Down	C--N--

VPN Flags: (E) Admin Enabled, (A) Oper Active, (I) Include Tag,
(X) Exclude Tag, (T) Ethertype Configured,


```
(V) VCCV HC Enabled, (W) VPN Type VPWS, (L) VPN Type VPLS
Peer Flags: (C) Core Peer, (S) Spoke Peer, (A) Active Core,
(p) Configured Primary Core, (s) Configured Secondary Core,
(N) Named LSP Configured, (V) VCCV HC Capabilities Negotiated,
(F) VCCV HC Failed
```

```
-----
Total number of configured L2VPNs:      9
Total number of active L2VPNs:          3
Total number of configured PWs:          4
Total number of active PWs:              3
PWs auto-selecting transport LSP:         1
PWs configured with a transport LSP:      3
PWs using LDP for transport:              0
PWs using RSVP for transport:            4
PWs using static for transport:           0
```

The following shows summary Layer 2 VPN information for VPLS peer 2.2.2.2:

```
# show l2vpn vpls peer 2.2.2.2
```

```
L2VPN Name      VPN ID  Flags  Services Name  Peer IP      State  Flags
-----
vsl             105    EAX--L cust1   2.2.2.2      UP        CAp-V-
VPN Flags: (E) Admin Enabled, (A) Oper Active, (I) Include Tag,
(X) Exclude Tag, (T) Ethertype Configured,
(V) VCCV HC Enabled, (W) VPN Type VPWS, (L) VPN Type VPLS
Peer Flags: (C) Core Peer, (S) Spoke Peer, (A) Active Core,
(p) Configured Primary Core, (s) Configured Secondary Core,
(N) Named LSP Configured, (V) VCCV HC Capabilities Negotiated,
(F) VCCV HC Failed
```

The following shows detailed Layer 2 VPN information for VPLS peer 11.100.100.210:

```
# show l2vpn vpls peer 11.100.100.210 detail
```

```
VPLS Name : vpls10
```

```
VPN ID           : 10                Admin State      : Enabled
Source Address    : 11.100.100.212   Oper State       : Enabled
VCCV Status       : Disabled          MTU              : 1500
VCCV Interval Time : 5 sec.           Ethertype        : 0x8100
VCCV Fault Multiplier : 4             .1q tag         : exclude
L2VPN Type        : VPLS              Redundancy       : None
Service Interface : vlan10
```

```
Peer IP : 11.100.100.210
```

```
PW State          : Up
PW Uptime         : 18d:0h:28m:26s
PW Installed      : True
Local PW Status   : No Faults
Remote PW Status  : No Faults
PW Mode           : Core-to-Core
Transport LSP     : LDP LSP (Not Configured)
Next Hop I/F      : o6vlan1
Next Hop Addr     : 12.182.0.216      Tx Label        : 0x00010
PW Rx Label       : 0x80405           PW Tx Label     : 0x80401
PW Rx Pkts        : 3806161633        PW Tx Pkts      : 4294967296
PW Rx Bytes       : 912385942         PW Tx Bytes     : 4294967296
MAC Limit         : No Limit
VCCV HC Status    : Not Sending (VCCV Not Enabled For This VPLS)
CC Type           : Rtr Alert          Total Pkts Sent : 0
CV Type           : LSP Ping           Total Pkts Rcvd : 0
```

```

Send Next Pkt      : --

Total Failures    : 0          Pkts During Last Failure : 0
Last Failure Tm   : --

```

The detail version of this command now displays a “PW Signaling” line that will display “LDP” or “None (Static)”, depending on the PW configuration. The “Local PW Status” will show “--” instead of “Not Signaled”, since the PW status is not currently signaled; however, for informational purposes, any local faults are still shown. The “Remote PW Status” and “Remote I/F MTU” will always show “--”.

Since the configured labels can be changed while the current labels are in-use, there is a small window where the configured labels and in-use labels are different. If you issue the `show l2vpn detail` command during this window, an extra line will be output to indicate this extra information. The configured labels are noted as “pending” in this case.

```

#show l2vpn vpls red04 detail
L2VPN Name: red04
  VPN ID          : 104          Admin State      : Enabled
  Source Address   : 11.100.100.102 Oper State      : Enabled
  VCCV Status      : Disabled    MTU              : 1500
  VCCV Interval Time : 5 sec.    Ethertype       : 0x8100
  VCCV Fault Multiplier : 4      .1q tag        : exclude
  L2VPN Type       : VPLS        Redundancy       : None
  Service Interface : red04svc

  Peer IP: 11.100.100.244
    PW State       : Up
    PW Signaling   : None (Static)
    PW Uptime      : 0d:0h:57m:32s
    PW Installed   : True
    Local PW Status : No Faults
    Remote PW Status : --
    Remote I/F MTU : --
    PW Mode        : Core-to-Core
    Transport LSP   : LDP LSP (Not Configured)
      Next Hop I/F  : e2-s4vlan1
      Next Hop Addr : 9.21.1.243      Tx Label      : 0x00248
    PW Rx Label    : 0x0002c          PW Tx Label   : 0x000c2
    PW Rx Label (pend) : 0x0002b      PW Tx Label (pend) : 0x000b2
    PW Rx Pkts     : 0                PW Tx Pkts    : 0
    PW Rx Bytes    : 0                PW Tx Bytes   : 0
    MAC Limit      : No Limit
    VCCV HC Status : Not Sending (VCCV Not Supported For Static PWs)
      CC Type      : --              Total Pkts Sent : 0
      CV Type      : --              Total Pkts Rcvd : 0
      Send Next Pkt : --
      Total Failures : 0              Pkts During Last Failure : 0
      Last Failure Tm : --

```

History

This command was first available in ExtremeXOS 11.6.

This command was updated to display flags for H-VPLS spoke nodes and protected VPLS and H-VPLS in ExtremeXOS 12.1.

The output for this command was modified in ExtremeXOS 12.2.2.

The **12vpn** and **vpws** keywords were first available in ExtremeXOS 12.4.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

show lacp

```
show lacp
```

Description

Displays LACP, or dynamic link aggregation, settings on the switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command displays the following information about the LACP [LAGs](#) configured on the switch:

- Up or Down.
- Enabled or disabled (not configurable).
- System MAC.
 - MAC address for the system, which is used for LACP priority in the absence of a specifically configured priority.
- LACP PDUs dropped on non-LACP ports.
- LAG.
 - Identifies the particular LAG. This number comes from logical port assigned to the LAG and is the LAG group ID.
- Actor Sys-Pri.
 - Shows the system priority for that LAG.
 - If this number is lower than the number displayed for the Partner Sys-Pri, the system you are working on is the controlling partner in the LAG.
- Actor Key.
 - Automatically generated LACP key.
- Partner MAC.
 - Identifies the MAC address for the system connecting to the LAG on the remote end.
- Partner Sys-Pri.
 - Shows the system priority for that LAG on the remote end.

- If this number is lower than the number displayed for the Actor Sys-Pri, the system at the remote end is the controlling partner in the LAG.
- Partner Key.
 - LACP key automatically generated by the system to which this aggregator is connected.
 - If this number is lower than the number displayed for the Actor Key, the partner system is the controlling partner in the LAG.
- Agg Count.
 - Identifies the number of ports added to the aggregator for that LAG.

Example

The following command displays the LACP LAGs on the switch:

```
show lacp
```

The following is sample output from this command on a switch:

```
LACP Up : Yes
LACP Enabled : Yes
System MAC : 00:04:96:10:33:60
LACP PDUs dropped on non-LACP ports : 0
Lag      Actor      Actor      Partner      Partner  Partner  Agg
Sys-Pri  Key        MAC        Sys-Pri      Key        Count
-----
1         90         0x07d1     00:01:30:f9:9c:30  601       0x1391    2
5         100        0x0fa5     00:01:30:f9:9c:30  321       0x1f47    16
9         677        0x0fa9     00:01:30:f9:9c:30  87        0x0fa9    8
```

History

This command was first available in ExtremeXOS 11.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show lacp counters

```
show lacp counters
```

Description

Displays all LACP, or dynamic link aggregation, counters for all member ports in the system.

Syntax Description

This command has no parameters or variables.

Default

N/A.

Usage Guidelines

This command displays the following information for all link aggregation groups (LAGs):

- LACP PDUs dropped on non-LACP ports.
- LACP bulk checkpointed messages sent.
- LACP bulk checkpointed messages received.
- LACP PDUs checkpointed sent.
- LACP PDUs checkpointed received.
- LAG group ID.
- Member port.
- Packets received.
- Packets dropped from PDU error.
- Packets dropped because LACP is not enabled on this port.
- Packets dropped because sender's system MAC address matches that of receiver.
- Packets successfully transmitted.
- Packets with errors during transmission.

Example

The following command displays LACP counters:

```
show lacp counters
```

The following is sample output from this command on a switch:

```
LACP PDUs dropped on non-LACP ports : 519392
LACP Bulk checkpointed msgs sent    : 1
LACP Bulk checkpointed msgs recv    : 0
LACP PDUs checkpointed sent        : 575616
LACP PDUs checkpointed recv        : 0
```

Lag Group	Member Port	Rx Ok	Rx Drop PDU Err	Rx Drop Not Up	Rx Drop Same MAC	Tx Sent	Tx Ok	Tx Xmit Err
1:1	1:1	2169	0	0	0	2170	0	
1:2	2169	0	0	0	2170	0		
1:3	2169	0	0	0	2170	0		
1:4	2169	0	0	0	2170	0		
1:5	2169	0	0	0	2170	0		
1:6	2169	0	0	0	2170	0		
1:7	2169	0	0	0	2170	0		
1:8	2168	0	0	0	2169	0		

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show lacp lag

```
show lacp lag group-id {detail}
```

Description

Displays LACP, or dynamic link aggregation, settings for the specified LAG.

Syntax Description

<i>group-id</i>	Specifies the LAG group ID you want to display. This is the number of the port you configured as the logical port of the LAG.
detail	Show detailed information.

Default

N/A.

Usage Guidelines

This command displays the following information about the specified LACP LAG:

- LAG
 - Identifies the particular LAG. This number comes from logical port assigned to the LAG and is the LAG group ID.
- Actor Sys-Pri
 - Shows the system priority for that LAG.
 - If this number is lower than the number displayed for the Partner Sys-Pri, the system you are working on is the controlling partner in the LAG.
- Actor Key
 - Automatically generated LACP key.
- Partner MAC
 - Identifies the MAC address for the system connecting to the LAG on the remote end.
- Partner Sys-Pri
 - Shows the system priority for that LAG on the remote end.
 - If this number is lower than the number displayed for the Actor Sys-Pri, the system at the remote end is the controlling partner in the LAG.
- Partner Key
 - LACP key automatically generated by the system to which this aggregator is connected.
 - If this number is lower than the number displayed for the Actor Key, the partner system is the controlling partner in the LAG.
- Agg Count
 - Identifies the number of ports added to the aggregator for that LAG.
- Member port
- Port priority

- Rx State—Receiving state of the port
 - Idle.
 - Initialized.
 - Current—Receiving LACP PDUs.
 - Expired.
 - Defaulted.
- Sel Logic—Selection state of the port
 - Selected—Ports with a matching admin key on the remote end.
 - Unselected—Ports that failed to meet with a matching admin key on the remote end.
 - Standby—Ports that exceed the number of ports that can be active in the LAG simultaneously. These ports can be moved into selected mode if one of the currently selected ports in the LAG goes down.
- Mux State—Ability to transmit and collect data of the port.
 - Waiting—Selected port that is waiting for LACP to determine if it can join the aggregator.
 - Attached—Ports ready to be added to the aggregator.
 - Collecting-Dist—Ports that are added to the aggregator and are transferring data.
 - Detached—Ports that cannot be added to the aggregator.
- Actor Flag—Mux state of the port.
 - A—Activity.
 - T—Timeout.
 - G—Aggregation.
 - S—Synchronization.
 - C—Collecting.
 - D—Distributing.
 - F—Defaulted.
 - E—Expired.
- Partner Port
 - The operational value of the port number assigned to this link by partner.
- Up—Yes or no.
- Enabled—Yes or no.
- LAG State—Up or Down.
- Unack count.
- Wait-for-count.
- Current timeout.
- Activity mode.
- Defaulted action.
- Receive state.
- Transmit state.
- Minimum Active—The minimum number of active links that must be up for the trunk to remain up.
- Selected count—Number of selected ports in the LAG.

- Standby count—Number of standby ports in the LAG.
- LAG Id flag
 - S—Displays information on controlling partner of LAG.
 - T—Displays information on controlled partner of LAG.

Example

The following command displays information on the specified LACP LAG:

```
show lacp lag 4:9
```

The following is sample output from this command on a switch:

```

Lag          Actor   Actor   Partner          Partner  Partner  Agg   Actor
      Sys-Pri Key     MAC              Sys-Pri  Key     Count Mac
-----
4:9          2110    0x0fa9  00:04:96:10:33:60  2110     0x0fa9   16    00:22:33:44:55:66

Port list:

Member      Port      Rx      Sel      Mux      Actor      Partner
Port        Priority  State   Logic    State     Flags      Port
-----
4:9         300      Current Selected Collect-Dist A-GSCD-- 4009
4:10        301      Current Selected Collect-Dist A-GSCD-- 4010
4:11        302      Current Standby   Detached  A-G----- 4011
4:12        303      Current Standby   Detached  A-G----- 4012
4:29        200      Current Selected Collect-Dist A-GSCD-- 4029
4:30         0       Current Selected Collect-Dist A-GSCD-- 4030
4:31        202      Current Selected Collect-Dist A-GSCD-- 4031
4:32        203      Current Selected Collect-Dist A-GSCD-- 4032
8:7         101      Current Selected Collect-Dist A-GSCD-- 8013
8:8         10       Current Selected Collect-Dist A-GSCD-- 8014
8:9         9        Current Selected Collect-Dist A-GSCD-- 8015
8:10        8        Current Selected Collect-Dist A-GSCD-- 8016
8:11        7        Current Selected Collect-Dist A-GSCD-- 8017
8:12        6        Current Selected Collect-Dist A-GSCD-- 8018
8:13        5        Current Selected Collect-Dist A-GSCD-- 8019
8:14        3        Current Selected Collect-Dist A-GSCD-- 8020
8:15        0        Current Selected Collect-Dist A-GSCD-- 8043
8:16        3        Current Selected Collect-Dist A-GSCD-- 8044
8:17        2        Idle    Unselected Detached  ----- 0
8:18        37       Idle    Unselected Detached  ----- 0
8:19        36       Idle    Unselected Detached  ----- 0
8:20        35       Idle    Unselected Detached  ----- 0
=====
Actor Flags: A-Activity, T-Timeout, G-Aggregation, S-Synchronization
              C-Collecting, D-Distributing, F-Defaulted, E-Expired

```

The following command displays detailed information on the specified LACP LAG:

```
show lacp lag 5 detail
```

The following is sample output from this command:

```

Lag   Actor   Actor   Partner          Partner  Partner  Agg   Actor
      Sys-Pri Key     MAC              Sys-Pri  Key     Count Mac

```



```

-----
5          0  0x03ed 00:04:96:52:5b:a3          0  0x03ed          4  00:04:96:52:5b:98

Enabled           : Yes
LAG State         : Down
Unack count       : 0
Wait-for-count    : 0
Current timeout   : Long
Activity mode     : Active
Defaulted Action  : Delete
Receive state     : Enabled
Transmit state    : Enabled
Minimum active    : <3
Selected count    : 2
Standby count     : 0
LAG Id flag       : Yes
  S.pri:0        , S.id:00:04:96:52:5b:98, K:0x03ed
  T.pri:0        , T.id:00:04:96:52:5b:a3, L:0x03ed

Port list:

Member   Port    Rx      Sel      Mux      Actor      Partner
Port     Priority State    Logic    State    Flags      Port
-----
5         0      Current  Selected  Attached  A-GSCD--  1005
6         0      Current  Selected  Attached  A-GSCD--  1006
7         0      Current  Unselected Attached  A-----  1007
8         0      Current  Unselected Attached  A-----  1008
=====
Actor Flags: A-Activity, T-Timeout, G-Aggregation, S-Synchronization
              C-Collecting, D-Distributing, F-Defaulted, E-Expired

```

History

This command was first available in ExtremeXOS 11.3.

Support for the LAG State and Minimum active parameters was added in ExtremeXOS 15.7.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show lacp member-port

```
show lacp member-port port {detail}
```

Description

Displays LACP, or dynamic link aggregation, settings for the specified port that is a member of any [LAG](#).

Syntax Description

<i>port</i>	Specifies the port number.
detail	Show detailed information.

Default

N/A.

Usage Guidelines

This command displays the following information about the specified port:

- Member Port.
- Port Priority.
- Rx State—Receiving state of the port.
 - Idle.
 - Initialized.
 - Current—Receiving LACP PDUs.
 - Expired.
 - Defaulted.
- Sel Logic—Selection state of the port.
 - Selected—Ports with a matching admin key on the remote end.
 - Unselected—Ports that failed to meet with a matching admin key on the remote end.
 - Standby—Ports that exceed the number of ports that can be active in the LAG simultaneously. These ports can be moved into selected mode if one of the currently selected ports in the LAG goes down.
- Mux State—Ability to transmit and collect data of the port.
 - Waiting—Selected port that is waiting for LACP to determine if it can join the aggregator.
 - Attached—Ports ready to be added to the aggregator.
 - Collecting-Dist—Ports that are added to the aggregator and are transferring data.
 - Detached—Ports that cannot be added to the aggregator.
- Actor Flag.
 - A—Activity.
 - T—Timeout.
 - G—Aggregation.
 - S—Synchronization.
 - C—Collecting.
 - D—Distributing.
 - F—Defaulted.
 - E—Expired.
- Partner Port.
 - The operational value of the port number assigned to this link by partner.
- Up or Down—LACP protocol running or not on specified port.

- Enabled or disabled (not configurable).
- Link State—Link state on this port up or down.
- Actor Churn—True or false.
- Partner Churn—True or false.
- Ready_N—Ready to be added to aggregator.
- Wait pending.
- Ack pending.
- LAG Id.
 - S—Displays information on controlling partner of LAG.
 - T—Displays information on controlled partner of LAG.
- Stats.
 - Rx - Accepted.
 - Rx - Dropped due to error in verifying PDU.
 - Rx - Dropped due to LACP not being up on this port.
 - Rx - Dropped due to matching own MAC.
 - Tx - Sent Successfully.
 - Tx - Transmit error.

Example

The following command displays LACP information on the specified port:

```
show lacp member-port 4:9
```

The following is sample output from this command on a switch:

Member Port	Port Priority	Rx State	Sel Logic	Mux State	Actor Flags	Partner Port
4:9	300	Current	Selected	Collect-Dist	A-GSCD--	4009

=====

Actor Flags: A-Activity, T-Timeout, G-Aggregation, S-Synchronization
C-Collecting, D-Distributing, F-Defaulted, E-Expired

The following command displays detailed LACP information on the specified port:

```
show lacp member-port 4:9 detail
```

The following is sample output from this command on a switch:

Member Port	Port Priority	Rx State	Sel Logic	Mux State	Actor Flags	Partner Port
4:9	300	Current	Selected	Collect-Dist	A-GSCD--	4009

Up : Yes
Enabled : Yes
Link State : Up
Actor Churn : False
Partner Churn : False
Ready_N : Yes

```
Wait pending : No
Ack pending  : No
LAG Id:
S.pri:2110, S.id:00:01:30:f9:9c:30, K:0x0fa9, P.pri:300 , P.num:4009
T.pri:2110, T.id:00:04:96:10:33:60, L:0x0fa9, Q.pri:300 , Q.num:4009
Stats:
Rx - Accepted : 2174
Rx - Dropped due to error in verifying PDU : 0
Rx - Dropped due to LACP not being up on this port : 0
Rx - Dropped due to matching own MAC : 0
Tx - Sent successfully : 2175
Tx - Transmit error : 0
=====
Actor Flags: A-Activity, T-Timeout, G-Aggregation, S-Synchronization
C-Collecting, D-Distributing, F-Defaulted, E-Expired
```

History

This command was first available in ExtremeXOS 11.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show ldap domain

```
show ldap domain {domain_name | all}
```

Description

This command displays the LDAP servers and other LDAP configuration details of one or all LDAP domains.

Syntax Description

<i>domain_name</i>	Displays the details of the specified domain.
all	Displays the details for all domains.

Default

N/A.

Usage Guidelines

Use this command to display the LDAP servers and other LDAP configuration details of one or all LDAP domains. The summary version (show ldap domain) displays the list of LDAP domains configured.

Example

```
# show ldap domain
```

```

-----
LDAP Domains
-----
XYZCorp.com (Default)
engg.XYZCorp.com
mktg.XYZCorp.com
sales.XYZCorp.com
-----

```

If no default domain is configured, this note appears at the bottom:

```

Note: No default domain configured
# show ldap domain all
-----
Domain(default) : XYZCorp.com
-----
Base-DN          : XYZCorp.com
Bind credential  : jsmith
LDAP Hierarchical Search OID : LDAP_MATCHING_RULE_IN_CHAIN (1.2.840.113556.1.4.1941)
LDAP Configuration for Netlogin:
dot1x           : Enabled
mac             : Enabled
web-based       : Enabled
LDAP Server 1   : 192.168.2.101
Server Port     : 389
Client IP       : Any
Client VR       : VR-Mgmt
Security Mechanism : Plain Text
Status          : Active
LDAP Server 2   : 192.168.2.102
Server Port     : 389
Client IP       : Any
Client VR       : VR-Mgmt
Security Mechanism : Plain Text
Status          : Not Active
-----
Domain          : engg.XYZCorp.com
-----
Base-DN          : engg.XYZCorp.com
Bind credential  : pkumar
LDAP Hierarchical Search OID : 1.2.840.113345.1.4.1789
LDAP Configuration for Netlogin:
dot1x           : Enabled
mac             : Enabled
web-based       : Enabled
LDAP Server 1   : engsrv1.engg.XYZCorp.com(192.168.3.101)
Server Port     : 389
Client IP       : 192.168.10.31
Client VR       : VR-Mgmt
Security Mechanism : Plain Text
Status          : Active
LDAP Server 2   : 192.168.3.102
Server Port     : 389
Client IP       : 192.168.10.31
Client VR       : VR-Mgmt
Security Mechanism : Plain Text
Status          : Not Active
-----
Domain          : it.XYZCorp.com
-----
Base-DN          : it.XYZCorp.com
Bind credential  : asingh
LDAP Hierarchical Search OID : None
LDAP Configuration for Netlogin:

```

```

dot1x      : Enabled
mac        : Enabled
web-based  : Enabled
LDAP Server 1      : 192.168.4.101
Server Port      : 389
Client IP       : 192.168.10.31
Client VR      : VR-Mgmt
Security Mechanism : Plain Text
Status         : Not Active
LDAP Server 2      : 192.168.4.102
Server Port      : 389
Client IP       : 192.168.10.31
Client VR      : VR-Mgmt
Security Mechanism : Plain Text
Status         : Active
-----
Domain      : mktg.XYZCorp.com
-----
Base-DN     : mktg.XYZCorp.com
Bind credential : gprasad
LDAP Hierarchical Search OID : LDAP_MATCHING_RULE_IN_CHAIN (1.2.840.113556.1.4.1941)
LDAP Configuration for Netlogin:
dot1x      : Enabled
mac        : Enabled
web-based  : Enabled
LDAP Server 1      : mktgsrv1.mktg.XYZCorp.com(192.168.5.101)
Server Port      : 389
Client IP       : Any
Client VR      : VR-Mgmt
Security Mechanism : Plain Text
Status         : Active
LDAP Server 2      : 192.168.5.102
Server Port      : 389
Client IP       : Any
Client VR      : VR-Mgmt
Security Mechanism : Plain Text
Status         : Not Active
-----
Domain      : sales.XYZCorp.com
-----
Base-DN     : sales.XYZCorp.com
Bind credential : masiq
LDAP Hierarchical Search OID : LDAP_MATCHING_RULE_IN_CHAIN (1.2.840.113556.1.4.1941)
LDAP Configuration for Netlogin:
dot1x      : Enabled
mac        : Enabled
web-based  : Enabled
LDAP Server  : No LDAP Servers configured
# show ldap domain "engg.XYZCorp.com"
-----
Domain      : engg.XYZCorp.com
-----
Base-DN     : engg.XYZCorp.com
Bind credential : pkumar
LDAP Hierarchical Search OID : LDAP_MATCHING_RULE_IN_CHAIN (1.2.840.113556.1.4.1941)
LDAP Configuration for Netlogin:
dot1x      : Enabled
mac        : Enabled
web-based  : Enabled
LDAP Server 1      : engsrv1.engg.XYZCorp.com(192.168.3.101)
Server Port      : 389
Client IP       : 192.168.10.31
Client VR      : VR-Mgmt
Security Mechanism : Plain Text

```

```
Status           : Active
LDAP Server 2    : 192.168.3.102
Server Port     : 389
Client IP       : 192.168.10.31
Client VR       : VR-Mgmt
Security Mechanism : Plain Text
Status         : Not Active
```

If the server was specified as a host name and the IP address was not resolved, this is shown:

```
LDAP Server1    : server1.domain.com(IP address unresolved)
```

History

This command was first available in ExtremeXOS 15.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show ldap statistics

```
show ldap statistics
```

Description

This command displays LDAP packet statistics per LDAP domain.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to show all LDAP related statistics per LDAP domain.

Example

```
Switch.21 # show ldap statistics
-----
Domain           : XYZCorp.com (default)
-----
LDAP Server 1    : 192.168.2.101
Server Port     : 389
Client VR       : VR-Mgmt
Status         : Active
```

```

Requests      : 12
Responses     : 12
Errors        : 0
LDAP Server 2 : 192.168.2.102
Server Port   : 389
Client VR     : VR-Mgmt
Status        : Not Active
Requests      : 0
Responses     : 0
Errors        : 0
-----
Domain        : engg.XYZCorp.com
-----
LDAP Server 1 : engsrv1.engg.XYZCorp.com(192.168.3.101)
Server Port   : 389
Client VR     : VR-Mgmt
Status        : Active
Requests      : 22
Responses     : 20
Errors        : 2
LDAP Server 2 : 192.168.3.102
Server Port   : 389
Client VR     : VR-Mgmt
Status        : Not Active
Requests      : 0
Responses     : 0
Errors        : 0
-----
Domain        : it.XYZCorp.com
-----
LDAP Server 1 : 192.168.4.101
Server Port   : 389
Client VR     : VR-Mgmt
Status        : Not Active
Requests      : 1
Responses     : 0
Errors        : 1
LDAP Server 2 : 192.168.4.102
Server Port   : 389
Client VR     : VR-Mgmt
Status        : Active
Requests      : 6
Responses     : 6
Errors        : 0
-----
Domain        : mktg.XYZCorp.com
-----
LDAP Server 1 : 192.168.5.101
Server Port   : 389
Client VR     : VR-Mgmt
Status        : Not Active
Requests      : 8
Responses     : 7
Errors        : 1
LDAP Server 2 : 192.168.5.102
Server Port   : 389
Client VR     : VR-Mgmt
Status        : Active
Requests      : 12
Responses     : 12
Errors        : 0
-----
Domain        : sales.XYZCorp.com

```



```
-----  
LDAP Server      : No LDAP Servers configured
```

History

This command was first available in ExtremeXOS 15.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show licenses

```
show licenses {[slot slot |all]} {detail}
```

Description

Displays current software license level, port speed licensing, and feature packs enabled on your switches.

Syntax Description

all	Specifies all slots on a stack.
<i>slot</i>	Specifies one or more slots on a stack.
detail	Specifies showing detailed license information. ExtremeSwitching Universal switches only.

Default

N/A.

Usage Guidelines

The command displays information on the software license level and feature packs enabled on the switch or nodes on a stack, including the trial license or factory default license and days left to expiry.

To see the license information on a SummitStack node, run the command while logged into that node.



Note

See the specific chapter of the [Switch Engine 32.3 User Guide](#) that discusses each feature to determine if a license is required for some functionality. If not noted, all functionality is available, and license is not required.

For ExtremeSwitching Universal switches, the DEF-EVAL-PRMR is a 30-day evaluation license that is enabled on the switch from the factory, or after changing the network operating to Switch Engine from Fabric Engine, or after a rescue boot.

Example

The following command displays the license level configuration:

```
show licenses
```

On a SummitStack, the output from this command looks similar to the following:

```
Slot-3 Stack.12 > show licenses
Enabled License Level:
Advanced Edge
Enabled Feature Packs:
None
Effective License Level:Edge
```

The show output also includes the current status of the non-Extreme optical device license:

```
# show license
Enabled License Level:
    Advanced Core
Enabled Feature Packs:
    MPLS ServiceProviderEdge 3rdPartyOptics
```

To see all license levels on a stack:

```
Slot-1 Stack.6 # show licenses slot all
```

Slot	Enabled License	Effective License	Feature Packs
1	Advanced Edge	Advanced Edge	DirectAttach
2	Advanced Edge	Advanced Edge	ServiceProviderEdge
			DirectAttach
			NetworkTiming

To view detailed license information on the ExtremeSwitching Universal switches:

```
# show licenses detail
```

Slot	License Product	Feature(s)	Installed	Effective	Expiration
1	DEF-EVAL-PRMR	Premier	Yes	No	Mon Sep 14 12:41:22 2020
1	PRD-5000-MACSEC	MACsec	No	N/A	N/A
2	PRD-5000-PRMR	Premier	No	N/A	N/A
2	PRD-5000-MACSEC	MACsec	Yes	N/A	Never
3	PRD-5000-PRMR	Premier	No	N/A	N/A
3	PRD-5000-MACSEC	MACsec	No	N/A	N/A

History

This command was first available in ExtremeXOS 11.1.

The information on enabled feature packs was added in ExtremeXOS 11.4.

The information on the trial licenses was added in ExtremeXOS 11.6.

The show output was updated to include the current status of the non-Extreme optical device license in ExtremeXOS 15.4.

The options **all** and **slot** were added in ExtremeXOS 22.5.

The **detail** option and ExtremeSwitching unified license information was added in ExtremeXOS 31.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show lldp

```
show lldp {port [all | port_list]} {detailed}
```

Description

Displays *LLDP* configuration information for the specified port or ports.

Use the detailed keyword to display the configured *VLANs* on the port and the enabled VLAN-specific TLVs.

Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.
detailed	Shows information on the configured VLANs on the port.

Default

N/A.

Usage Guidelines

Use the detailed variable to display information regarding configured VLANs on the ports and any enabled VLAN-specific TLVs.

Example

The following example displays LLDP configuration information for the switch:

```
# show lldp
LLDP transmit interval      : 30 seconds
LLDP transmit hold multiplier : 4  (used TTL = 120 seconds)
LLDP transmit delay        : 2 seconds
LLDP SNMP notification interval : 5 seconds
LLDP reinitialize delay     : 2 seconds
LLDP-MED fast start repeat count : 4
LLDP local management address : VLAN Chicago (10.127.2.30)
LLDP Port Configuration:
Port    Rx      Tx      SNMP      --- Optional enabled transmit TLVs ---
Mode    Mode    Notification  LLDP    802.1  802.3  MED    AvEx  DCBX
=====
1       Enabled Enabled  --      --D--  ---  ----  ----  IB
2       Enabled Enabled  --      --D--  ---  ----  ----  IB
=====
Notification: (L) lldpRemTablesChange, (M) lldpXMedTopologyChangeDetected
LLDP Flags   : (P) Port Description, (N) System Name, (D) System Description
              (C) System Capabilities, (M) Mgmt Address
```

```

802.1 Flags : (P) Port VLAN ID, (p) Port & Protocol VLAN ID, (N) VLAN Name
802.3 Flags : (M) MAC/PHY Configuration/Status, (P) Power via MDI
(+) Power via MDI with DLL Classification for PoE+,
(L) Link Aggregation, (F) Frame Size
MED Flags   : (C) MED Capabilities, (P) Network Policy,
(L) Location Identification, (p) Extended Power-via-MDI
AvEx Flags  : (P) PoE Conservation Request, (C) Call Server, (F) File Server
(Q) 802.1Q Framing
DCBX Flags  : (I) IEEE 802.1Qaz DCBX, (B) Baseline v1.01 DCBX

```

The following example includes detailed information on the LLDP configuration for port 1:1:

```

# show lldp port 1:1 detailed
LLDP transmit interval      : 30 seconds
LLDP transmit hold multiplier : 4  (used TTL = 120 seconds)
LLDP transmit delay         : 2 seconds
LLDP SNMP notification interval : 5 seconds
LLDP reinitialize delay     : 2 seconds
LLDP-MED fast start repeat count : 4
LLDP local management address : VLAN Chicago (10.127.2.30)
LLDP Port Configuration:
Port    Rx      Tx      SNMP      --- Optional enabled transmit TLVs ---
Mode    Mode    Notification  LLDP    802.1  802.3  MED    AvEx    DCBX
=====
1:1     Enabled  Enabled    --      --D--  ---    ----    CLP-  ----  IB
VLAN: Default
VLAN: voice
AvEx Call-Server: IP Address(es)=10.0.0.20, 10.0.0.21
AvEx File-Server: IP Address(es)=10.0.0.20, 10.0.0.21, 10.0.0.22
AvEx 802.1Q Framing: Mode=tagged
MED LCI: Location Format=ECS ELIN based
1234567890
MED Policy: Application=voice
VLAN=voice, DSCP=40
DCBX: Priority 4, iSCSI
DCBX: Priority 3, FCoE
DCBX: Priority 3, FIP
=====
Notification: (I) lldpRemTablesChange, (M) lldpXMedTopologyChangeDetected
LLDP Flags   : (P) Port Description, (N) System Name, (D) System Description
(C) System Capabilities, (M) Mgmt Address
802.1 Flags : (P) Port VLAN ID, (p) Port & Protocol VLAN ID, (N) VLAN Name
802.3 Flags : (M) MAC/PHY Configuration/Status, (P) Power via MDI
(L) Link Aggregation, (F) Frame Size
MED Flags   : (C) MED Capabilities, (P) Network Policy,
(L) Location Identification, (p) Extended Power-via-MDI
AvEx Flags  : (P) PoE Conservation Request, (C) Call Server, (F) File Server
(Q) 802.1Q Framing
DCBX Flags  : (I) IEEE 802.1Qaz DCBX, (B) Baseline v1.01 DCBX

```

History

This command was first available in ExtremeXOS 11.2.

The information on fast-start repeat count, MED, AvEx, and notification was added in ExtremeXOS 11.5.

An additional flag was added for PoE+ in ExtremeXOS 12.5.

The display was updated for DCBX in ExtremeXOS 12.6.

Management IP address information was added in ExtremeXOS 22.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show lldp dcbx

```
show lldp {port [all | port_list]} dcbx {ieee|baseline} {detailed}
```

Description

Displays DCBX configuration and statistics information for one or all ports.

Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.
ieee	Specifies IEEE 802.1Qaz information only.
baseline	Specifies Baseline v1.01 information only.
detailed	Shows information on the configured VLANs on the port.

Default

N/A.

Usage Guidelines

The summary display (without the detailed option) displays the status for each DCBX TLV on each port. For each TLV, the status is reported as shown in the following table.

DISABLED	DCBX is disabled on the port. This port status appears only in the summary display when DCBX is enabled for one version and disabled for the other. In the detailed display, ports on which DCBX are disabled are not shown.
OK	This TLV has been received by the peer, and either the configuration matches, or the peer is reporting that it is in willing mode and is not reporting an explicit error.
UNKNOWN	This TLV has not been received by the peer since the port has been active.
EXPIRED	This TLV has been received by the peer, but the time to live has expired.
ERROR	Either a mismatch exists between the local and remote configuration and the peer is not willing, or the peer is reporting an error.
MULTIPLE PEERS	More than one <i>LLDP</i> peer has been detected on the link.

When you specify a port or the detailed option, local TLV information includes the information that will be contained in the next TLV that is sent, and if the configuration hasn't changed, this is the same information that was sent in the last TLV. Peer TLV information displays the information from the last TLV that has been received. For each TLV, statistics are reported as follows:

- Sent: Total number of TLVs sent since port has been operational.
- Received: Total number of TLVs received since port has been operational.
- Errors: Total number of mal-formed TLVs received since port has been operational.

You can clear the statistics using the clear counters command.

Table 27 describes the IEEE 802.1Qaz DCBX TLVs that can be displayed. Table 28 on page 2804 describes the Baseline v1.01 DCBX TLVs.

Table 27: IEEE 802.1Qaz DCBX TLVs

TLV/Description	Contents/Description
ETS TLV Advertises the ETS configuration of the local port and the configuration recommended to/by the peer for the specified port, respectively.	<p>Willing—Whether or not the device is willing to accept configuration from its DCBX peer. Zero (0) means No, and one (1) means Yes. The Willing bit does not apply to the ETS Recommendation TLV, and should always be zero.</p> <p>CBS—Whether the device supports the credit-based shaper algorithm. Zero (0) means No, and one (1) means Yes.</p> <p>Max TCs— Maximum number of traffic classes that the node can support.</p> <p>Priority Assgn—Priority Assignment Table. A priority group (PG) table describing how 802.1p priorities are assigned to PGs. The table is laid out as follows: Priority-0 : Priority-1 : Priority-2 : Priority-3 : Priority-4 : Priority-5 : Priority-6 : Priority-7 The value in the Priority-N position indicates the TC ID to which packets with an 802.1p priority of N are mapped.</p> <p>Note: For Extreme Networks products, a traffic class (TC) is synonymous with a <i>QoS Profile</i> (QP), except that TCs are zero-based, and QPs are one-based, so TC 1 maps to QP 0.</p> <p>TC Bwdth—TC Bandwidth Table. Indicates the percentage of bandwidth allocated for each traffic class. The table is laid out as follows: TC%-0 : TC%-1 : TC%-2 : TC%-3 : TC%-4 : TC%-5 : TC%-6 : TC%-7 The value in the TC%-N position indicates the percentage of the link bandwidth allocated to TC N. The total of all positions must equal 100.</p> <p>TSA—Transmission Selection Algorithm (TSA) Assignment Table. The table is laid out as follows: TC-0 : TC-1 : TC-2 : TC-3 : TC-4 : TC-5 : TC-6 : TC-7 The value in the TC-N position indicates the TSA used by TC N, which is one of the following: S - Strict priority (TSA 0)C - Credit-based shaper (TSA 1)E - Enhanced Transmission Selection (TSA 2)V - Vendor-specific Transmission Selection algorithm (TSA 255)</p> <p>Note: TSA values 3 to 254 are reserved for future standardization.</p>
Common Feature TLVs TLVs common to the Priority Group, PFC, and Application TLVs	<p>Oper Vers—Operating version of the feature.</p> <p>Max Vers—Highest feature version supported by the system.</p> <p>Enabled—Locally administered parameter that indicates whether the DCB feature is enabled. Zero (0) means No, and one (1) means Yes.</p> <p>Willing—Indicates whether the device is willing to accept configuration from its DCBX peer. Zero (0) means No, and one (1) means Yes.</p> <p>Error—Indicates whether an error has occurred during the configuration exchange with the peer. Zero (0) means No, and one (1) means Yes.</p>

Table 27: IEEE 802.1Qaz DCBX TLVs (continued)

TLV/Description	Contents/Description
Priority Group TLV Advertises priority to priority group mapping, priority group bandwidth and the scheduling algorithm.	<p>PG IDs—Priority Allocation Table. A priority group (PG) table describing how 802.1p priorities are assigned to PGs. The table is laid out as follows: Priority-0 : Priority-1 : Priority-2 : Priority-3 : Priority-4 : Priority-5 : Priority-6 : Priority-7 The value in the Priority-N position indicates the PG ID to which packets with an 802.1p priority of N are mapped. If the value is in the range of 0 to 7, this is the actual PG. If the value is equal to 15, this priority is mapped to a non-ETS group. In the case of Extreme Networks products, this would be a strict priority group.</p> <p>Note: For Extreme Networks products, a priority group (PG) is synonymous with a QoS Profile (QP), except that PGs are zero-based, and QPs are one-based, so PG1 maps to QP 0.</p> <p>PG%—Priority Group Allocation Table. Indicates the percentage of bandwidth allocated for each priority group. The table is laid out as follows: PG%-0 : PG%-1 : PG%-2 : PG%-3 : PG%-4 : PG%-5 : PG%-6 : PG%-7 The value in the PG%-N position indicates the percentage of the link bandwidth allocated to PG N. The total of all positions must equal 100.</p> <p>Num TCs—Maximum number of priority groups that the node can support.</p>
PFC TLV Describes the PFC configuration for the given port.	<p>Willing—Whether or not the device is willing to accept configuration from its DCBX peer. Zero (0) means No, and one (1) means Yes.</p> <p>MBC—MACsec Bypass Capability. If set to zero (0), the device is capable of bypassing MACsec processing when MACsec is disabled. If set to one (1), the sending station is not capable of bypassing MACsec processing when MACsec is disabled.</p> <p>PFC Cap—PFC Capability. The maximum number of classes on which the device may simultaneously support PFC.</p> <p>PFC Enable—List of priorities on which PFC is enabled.</p>
Application TLV Displays the priority the device expects to be used for the specified application.	<p>Priority—The priority to be used for the given protocol.</p> <p>Application—Specifies one of the following:</p> <ul style="list-style-type: none"> • FCoE • FIP • iSCSI • EtherType: <i>ethertype</i> • TCP/UDP Port: <i>port number</i> • TCP Port: <i>port number</i> • TCP Port: <i>port number</i>

Table 28: Baseline v1.01 DCBX TLVs

TLV/Description	Contents/Description
Control TLV Contains general information about the DCBX session.	<p>Oper Vers—Operating version of the DCBX protocol.</p> <p>Max Vers—Highest DCBX protocol version supported by the system.</p>

Table 28: Baseline v1.01 DCBX TLVs (continued)

TLV/Description	Contents/Description
	<p>Seq No—A value that changes each time an exchanged parameter in one or more of the DCB feature TLVs changes.</p> <p>Ack No—The SeqNo value from the most recent peer DCBX TLV that has been handled. This value acknowledges to the peer that a specific SeqNo has been received.</p>
Common Feature TLVs TLVs common to the Priority Group, PFC, and Application TLVs	<p>Oper Vers—Operating version of the feature.</p> <p>Max Vers—Highest feature version supported by the system.</p> <p>Enabled—Locally administered parameter that indicates whether the DCB feature is enabled. Zero (0) means No, and one (1) means Yes.</p> <p>Willing—Indicates whether the device is willing to accept configuration from its DCBX peer. Zero (0) means No, and one (1) means Yes.</p> <p>Error—Indicates whether an error has occurred during the configuration exchange with the peer. Zero (0) means No, and one (1) means Yes.</p>
Priority Group TLV Advertises priority to priority group mapping, priority group bandwidth and the scheduling algorithm.	<p>PG IDs—Priority Allocation Table. A priority group (PG) table describing how 802.1p priorities are assigned to PGs. The table is laid out as follows: Priority-0 : Priority-1 : Priority-2 : Priority-3 : Priority-4 : Priority-5 : Priority-6 : Priority-7 The value in the Priority-N position indicates the PG ID to which packets with an 802.1p priority of N are mapped. If the value is in the range of 0 to 7, this is the actual PG. If the value is equal to 15, this priority is mapped to a non-ETS group. In the case of Extreme Networks products, this would be a strict priority group.</p> <p>Note: For Extreme Networks products, a priority group (PG) is synonymous with a QoS Profile (QP), except that PGs are zero-based, and QPs are one-based, so PG1 maps to QP 0.</p> <p>PG%—Priority Group Allocation Table. Indicates the percentage of bandwidth allocated for each priority group. The table is laid out as follows: PG%-0 : PG%-1 : PG%-2 : PG%-3 : PG%-4 : PG%-5 : PG%-6 : PG%-7 The value in the PG%-N position indicates the percentage of the link bandwidth allocated to PG N. The total of all slots must equal 100.</p> <p>Num TCs—Maximum number of priority groups that the node can support.</p>
PFC TLV Describes the PFC configuration for the given port.	<p>PFC Enable—List of priorities on which PFC is enabled.</p> <p>Num TC PFCs—The maximum number of classes on which the device may simultaneously support PFC.</p>
Application TLV Displays the priority the device expects to be used for the specified application.	<p>Priority—The priority to be used for the given protocol.</p> <p>Application—Specifies one of the following:</p> <ul style="list-style-type: none"> • FCoE • FIP • iSCSI • EtherType: <i>ethertype</i> • TCP/UDP Port: <i>port number</i>

Example

The following example displays the summary DCBX configuration and statistics:

```
# show lldp dcbx
```

Baseline DCBX TLV Status:				IEEE DCBX TLV Status:				
Port	Control	PG	PFC	App	ETS-Conf	ETS-Rec	PFC	App
1	OK	OK	OK	OK	OK	OK	OK	OK
2	OK	OK	OK	OK	OK	OK	OK	OK
3	OK	OK	OK	OK	OK	OK	OK	OK
4	OK	OK	OK	OK	OK	OK	OK	OK
5	UNKNOWN	UNKNOWN	UNKNOWN	UNKNOWN	UNKNOWN	UNKNOWN	UNKNOWN	UNKNOWN
9	UNKNOWN	UNKNOWN	UNKNOWN	UNKNOWN	UNKNOWN	UNKNOWN	UNKNOWN	UNKNOWN
10	UNKNOWN	UNKNOWN	UNKNOWN	UNKNOWN	DISABLED	DISABLED	DISABLED	DISABLED
16	DISABLED	DISABLED	DISABLED	DISABLED	UNKNOWN	UNKNOWN	UNKNOWN	UNKNOWN
23	UNKNOWN	UNKNOWN	UNKNOWN	UNKNOWN	UNKNOWN	UNKNOWN	UNKNOWN	UNKNOWN
24	UNKNOWN	UNKNOWN	UNKNOWN	UNKNOWN	UNKNOWN	UNKNOWN	UNKNOWN	UNKNOWN

```

Control - Control TLV
PG       - Priority Group TLV
PFC      - Priority-Based Flow Control TLV
App      - Application Configuration TLV
ETS-Conf - ETS Configuration TLV
ETS-Rec  - ETS Recommendation TLV

```

The following example displays detailed IEEE 802.1Qaz DCBX configuration and statistics information for port 1:

```
# show lldp ports 1 dcbx ieee
Port number : 1
IEEE 802.1Qaz DCBX Information:
-----
ETS Configuration TLV: Sent: 5996, Received: 5997, Errors: 0, Status: OK
Local TLV : Willing: 0, CBS: 1, Max TCs: 8
Priority Assgn: 0:0:0:0:0:0:7, TC Bwdth: 33:0:0:33:34:0:0:0, TSA: E:S:S:E:E:S:S:S
Peer TLV : Willing: 0, CBS: 1, Max TCs: 8
Priority Assgn: 0:0:0:0:0:0:7, TC Bwdth: 33:0:0:33:34:0:0:0, TSA: E:S:S:E:E:S:S:S
ETS Recommendation TLV: Sent: 5996, Received: 5997, Errors: 0, Status: OK
Local TLV : Willing: 0, CBS: 0, Max TCs: 8
Priority Assgn: 0:0:0:0:0:0:7, TC Bwdth: 33:0:0:33:34:0:0:0, TSA: E:S:S:E:E:S:S:S
Peer TLV : Willing: 0, CBS: 0, Max TCs: 8
Priority Assgn: 0:0:0:0:0:0:7, TC Bwdth: 33:0:0:33:34:0:0:0, TSA: E:S:S:E:E:S:S:S
PFC TLV: Sent: 5996, Received: 5997, Errors: 0, Status: OK
Local TLV : Willing: 0, MBC: 0, Max PFCs: 8, PFC Enable: 3,4
Peer TLV : Willing: 0, MBC: 0, Max PFCs: 8, PFC Enable: 3,4
Application TLV: Sent: 5987, Received: 5988, Errors: 0, Status: OK
Local TLV : Priority: 4, iSCSI
Priority: 3, FCoE
Priority: 3, FIP
Peer TLV : Priority: 4, iSCSI
Priority: 3, FCoE
Priority: 3, FIP
```

The following example displays detailed Baseline v1.01 DCBX configuration and statistics information for port 1:

```
# show lldp ports 1 dcbx baseline
Port number : 1
Baseline v1.01 DCBX Information:
```

```
-----
Control TLV: Sent: 5999, Received: 6000, Errors: 0, Status: OK
Local TLV : Oper Vers: 0, Max Vers: 0, Seq No: 17, Ack No: 17
Peer TLV  : Oper Vers: 0, Max Vers: 0, Seq No: 17, Ack No: 17
Priority Group TLV: Sent: 5999, Received: 6000, Errors: 0, Status: OK
Local TLV : Oper Vers: 0, Max Vers: 0, Enabled: 1, Willing: 0 Error: 0
PG IDs: 0:0:0:0:0:0:15, PG%: 33:0:0:33:34:0:0:0, Num TCs: 8
Peer TLV  : Oper Vers: 0, Max Vers: 0, Enabled: 1, Willing: 0 Error: 0
PG IDs: 0:0:0:0:0:0:15, PG%: 33:0:0:33:34:0:0:0, Num TCs: 8
PFC TLV: Sent: 5999, Received: 6000, Errors: 0, Status: OK
Local TLV : Oper Vers: 0, Max Vers: 0, Enabled: 1, Willing: 0 Error: 0
Max PFCs: 8, PFC Enable: 3,4
Peer TLV  : Oper Vers: 0, Max Vers: 0, Enabled: 1, Willing: 0 Error: 0
Max PFCs: 8, PFC Enable: 3,4
App TLV: Sent: 5990, Received: 5991, Errors: 0, Status: OK
Local TLV : Oper Vers: 0, Max Vers: 0, Enabled: 1, Willing: 0 Error: 0
Priority: 4, iSCSI
Priority: 3, FCoE
Priority: 3, FIP
Peer TLV  : Oper Vers: 0, Max Vers: 0, Enabled: 1, Willing: 0 Error: 0
Priority: 4, iSCSI
Priority: 3, FCoE
Priority: 3, FIP
```

History

This command was first available in ExtremeXOS 12.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show lldp neighbors

```
show lldp {port [all | port_list]} neighbors {detailed}
```

Description

Displays the information related to the LLDP neighbors detected on the specified port or ports.

Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.
detailed	Shows detailed information on the neighbors.

Default

N/A.

You must use the detailed parameter to display detailed information about the received LLDP TLVs.

The following example displays LLDP neighbor information for all switch ports:

[illegible]

The following command lists detailed LLDP neighbor information for all switch ports:

```
# show lldp neighbors detailed
LLDP Port 17 detected 1 neighbor
Neighbor: 00:04:96:98:80:96/17, age 2 seconds
- Chassis ID type: MAC address (4)
  Chassis ID      : 00:04:96:98:80:96
- Port ID type: ifName (5)
  Port ID        : "17"
- Time To Live: 120 seconds
- System Name: "X670G2-48x-4q"
- System Description: "ExtremexOS (X670G2-48x-4q) version 30.2.0.24 xos_
                      30.2 by tlewis on Tue Jan 29 09:31:57 EST 2019"
- Avaya/Extreme Fabric Attach element
  Element Type    : 2
  State           : 0
  Management Vlan: 0
  SystemId        : 00:04:96:98:80:96
  Connection Type: MLT
  SMLT-ID         : 258
  MLT-ID          : 768
  EXOS Slot:Port  : 1:1
  Authenticated   : Yes

- Extreme Fabric Attach NSI(I-SID)/VLAN map
  Authenticated   : Yes
Status  VlanID  NSI(I-SID)
-----
      2      2222      11123
```

The following example shows Port Extension TLV information for bridge port extenders:

```
# show lldp port 1:2 neighbors detailed
LLDP Port 1:2 detected 1 neighbor
  Neighbor: 02:04:96:97:DD:94/1:2, age 22 seconds
    - Chassis ID type: MAC address (4)
      Chassis ID      : 02:04:96:97:DD:94
    - Port ID type: ifName (5)
      Port ID         : "1:2"
```

```

- Time To Live: 120 seconds
- System Name: "Stack"
- System Description: "ExtremeXOS (Stack) version 22.2.0.7 dot1br_22.2.0\
                      .7 by release-manager on Wed Sep 7 14:33:29 EDT 2016"
- IEEE802.3 Link-aggregation
  Capability/Status : Capable, Enabled (0x03)
  Aggregated Port ID: 0 (0x00000000)
- MED Model Name: "BPE-G24t."
- IEEE802.1BR Port Extension:
  Priority          : 255
  PE CSP Address: 00:00:01:48:00:1a
  PE Address       : 00:00:01:48:00:00

```

The following example shows BGP Config TLVs. This TLV allows SLX switches to communicate with ExtremeXOS's BGP Auto-peering feature:

```

# show lldp neighbors detailed

-----
LLDP Port 18 detected 1 neighbor
Neighbor: 00:04:96:97:D1:5C/18, age 28 seconds
- Chassis ID type: MAC address (4)
  Chassis ID      : 00:04:96:97:D1:5C
- Port ID type: ifName (5)
  Port ID         : "18"
- Time To Live: 120 seconds
- System Name: "X460G2-48t-10G4"
- System Description: "ExtremeXOS (X460G2-48t-10G4) version 30.2.0.11 xo\
                      s_30.2 by kmalone on Mon Oct 29 13:46:24 EDT 2018"
- EXTR Capabilities: "EasyBGP"
- EXTR EasyBGP:
  EXTR EasyBGP AS-NUM : 100
  EXTR EasyBGP RouterId: 1.1.0.100
  EXTR EasyBGP Address : fe80::204:96ff:fe97:d15c
- BGP Config:
  BGP Config Peer Address      : fe80::204:96ff:fe97:d15c
  BGP Config Peer Address AFI/SAFI: IPv6/Unicast
  BGP Config Local-AS         : 100
  BGP Config Identifier       : 1.1.0.100

```

History

This command was first available in ExtremeXOS 11.2. Information on the LLDP MED extension and Avaya-Extreme proprietary TLVs was added in ExtremeXOS 11.5.

Additional PoE+ information can appear when present in a Power via MDI TLV received from a neighbor starting in ExtremeXOS 12.5.

"System Name" was added to the output in ExtremeXOS 22.1.

Management IP address information was added in ExtremeXOS 22.4.

Port Extension TLV information for bridge port extenders (BPEs) was added in ExtremeXOS 22.5.

Fabric Attach authentication information and BGP Config TLVs (allows SLX switches to communicate with ExtremeXOS's BGP auto-peering feature) were added in ExtremeXOS 30.2.

Fabric Attach detailed link information was added in ExtremeXOS 31.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show lldp statistics

```
show lldp {port [all | port_list]} statistics
```

Description

Displays statistical counters related to the specified port or ports.

Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

N/A.

Usage Guidelines

The following counters are presented with the standard command (taken from the IEEE 802.1ab MIB definition):

- Last table change time: Last time an entry in the LLDP database was added, changed or deleted.
- Number of table inserts: The number of times the complete set of information advertised by a particular neighbor has been inserted into tables.
- Number of table deletes: The number of times the complete set of information advertised by a particular neighbor has been deleted from tables.
- Number of table drops: The number of times the complete set of information advertised by a particular neighbor could not be stored in memory because of insufficient resources.
- Number of table age outs: The number of times the complete set of information advertised by a particular neighbor has been deleted from tables because the information timeliness interval has expired.
- Tx Total: The number of LLDP frames transmitted by this switch on the indicated port.
- Tx Total Length Exceeded: The number of LLDP frames sent out on this port that could not hold all the information configured because the total frame length would exceed the maximum LDDPDU size of 1500 bytes.
- Rx Total: The number of valid LLDP frames received by this switch on the indicated port, while this LLDP agent is enabled.
- Rx Discarded: The number of LLDP frames received by this switch on the indicated port, and then discarded for any reason.
- Rx Errors: The number of invalid LLDP frames received by this switch on the indicated port, while this LLDP agent is enabled.

- TLVs Discarded: The number of LLDP TLVs discarded for any reason by this switch on the indicated port.
- TLVs Unrecognized: The number of LLDP TLVs received on the given port that are not recognized by the switch.

Example

The following example lists statistical counters for all ports on the switch:

```
# show lldp port all statistics
Last table change time   : Fri Dec 17 10:42:33 2004
Number of Table Inserts  : 3
Number of Table Deletes  : 0
Number of Table Drops    : 0
Number of Table Age Outs : 0
Port      Tx          Tx Length Rx Rx          Rx          TLVs          TLVs
Total     Exceeded TotalDiscarded  Errors    Discarded  Unrecogn.
=====
1:1       189         05654      0          0          0          0
2:2       188         0565      0          0          0          0
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show log

```
show log {messages [memory-buffer | nvram]} {events {event-condition
| event-component}} {severity severity {only}} {starting [date date
time time | date date | time time]} {ending [date date time time |
date date | time time]} {match regex} {chronological}
```

Description

Displays the current log messages.

Syntax Description

messages	Specifies the target location from which to display the log messages.
memory-buffer	Show messages stored in volatile memory (default).
nvr am	Show messages stored in NVRAM.
events	Show event messages.
<i>event-condition</i>	Specifies the event condition to display.
<i>event-component</i>	Specifies the event component to display.

<i>severity</i>	Specifies the minimum severity level to display (if the keyword only is omitted).
only	Specifies that only the specified severity level is to be displayed
starting	Show messages with timestamps equal to or greater than that specified
<i>date</i>	Specifies the date, where date is <i>month (1-12) / day (1-31) {/ year (yyyy)}</i> .
<i>time</i>	Specifies the time, where time is <i>hour (0-23) {: minute (0-59) {: seconds (0-59) {. hundredths}}}</i> .
ending	Show messages with timestamps equal to or less than that specified.
<i>regex</i>	Specifies a regular expression. Only messages that match the regular expression will be displayed.
chronological	Specifies displaying log messages in ascending chronological order (oldest to newest).

Default

The following defaults apply:

- messages—memory buffer.
- event—no restriction (displays user-specified event).
- severity—none (displays everything stored in the target).
- starting, ending—if not specified, no timestamp restriction.
- match—no restriction.
- chronological—if not specified, show messages in order from newest to oldest.

Usage Guidelines

Switch configuration and fault information is filtered and saved to target logs, in a memory buffer, and in NVRAM. Each entry in the log contains the following information:

- Timestamp—records the month and day of the event, along with the time (hours, minutes, seconds, and hundredths).
- Severity Level—indicates the urgency of a condition reported in the log. [Table 29](#) on page 2814Table describes the severity levels assigned to events.
- Component, Subcomponent, and Condition Name—describes the subsystem in the software that generates the event. This provides a good indication of where a fault might lie.
- Message—a description of the event occurrence. If the event was caused by a user, the user name is also provided.

This command displays the messages stored in either the internal memory buffer or in NVRAM. The messages shown can be limited by specifying a severity level, a time range, or a match expression. Messages stored in the target have already been filtered as events occurred, and specifying a severity or match expression on the show log command can only further limit the messages shown.

If the messages keyword is not present, the messages stored in the memory-buffer target are displayed. Otherwise, the messages stored in the specified target are displayed.

If the only keyword is present following the severity value, then only the events at that exact severity are included. Without the only keyword, events at that severity or more urgent are displayed. For example, severity warning implies critical, error, or warning, whereas severity warning only implies only warning.

Messages whose timestamps are equal or later than the starting time and are equal or earlier than the specified ending time will be shown if they also pass the severity requirements and match expression, if specified.

If a match phrase is specified, the formatted message must match the simple regular expression specified by match-expression for it to be shown.

A simple regular expression is a string of single characters including the dot character (.), which are optionally combined with quantifiers and constraints. A dot matches any single character while other characters match only themselves (case is significant). Quantifiers include the star character (*) that matches zero or more occurrences of the immediately preceding character or dot. Constraints include the caret character (^) that matches at the beginning of a message, and the currency character (\$) that matches at the end of a message. Bracket expressions are not supported. There are a number of sources available on the Internet and in various language references describing the operation of regular expressions.

If the chronological keyword is specified, messages are shown from oldest to newest; otherwise, messages are displayed newest to oldest.

Severity Level

The severity levels are critical, error, warning, notice, and info, plus three severity levels for extended debugging, debug-summary, debug-verbose, and debug-data. In log messages, the severity levels are shown by four letter abbreviations. The abbreviated forms are:

- Critical—Crit.
- Error—Erro.
- Warning—Warn.
- Notice—Noti.
- Info—Info.
- Debug-Summary—Summ.
- Debug-Verbose—Verb.
- Debug-Data—Data.

The three severity levels for extended debugging, debug-summary, debug-verbose, and debug-data, require that debug mode be enabled (which may cause a performance degradation). See the command `enable log debug-mode`. The following table describes the severity levels:

Table 29: Severity Levels Assigned by the Switch

Level	Description
Critical	A serious problem has been detected that is compromising the operation of the system and that the system cannot function as expected unless the situation is remedied. The switch may need to be reset.
Error	A problem has been detected that is interfering with the normal operation of the system and that the system is not functioning as expected.
Warning	An abnormal condition, not interfering with the normal operation of the system, has been detected that may indicate that the system or the network in general may not be functioning as expected.
Notice	A normal but significant condition has been detected, which signals that the system is functioning as expected.
Info (Informational)	A normal but potentially interesting condition has been detected, which signals that the system is functioning as expected and simply provides information or confirmation about the condition.
Debug-Summary	A condition has been detected that may interest a developer determining the reason underlying some system behavior.
Debug-Verbose	A condition has been detected that may interest a developer analyzing some system behavior at a more verbose level than provided by the debug summary information.
Debug-Data	A condition has been detected that may interest a developer inspecting the data underlying some system behavior.

Messages stored in NVRAM are in encoded format. To restore the ASCII text of a message, the version of ExtremeXOS loaded must be able to interpret the data written prior to reboot. When the encoded format for a particular message cannot be interpreted by the version of ExtremeXOS currently loaded, the messages are displayed in the following format:

```
03/21/2005 17:15:37.36 : NO MESSAGE DECODE; Missing component "epm" v24.2 DUMP-10: 00 14
C3 C1 00 11 00 1C 01 FF 00 08 65 70 6D 00 '.....epm.' DUMP-20: 08 FF 00 0C 00 18 00
02 65 70 6D 00 '.....epm.'
```

Log entries remain in the NVRAM log after a switch reboot. Issuing a clear log command does not remove these static entries. To remove log entries from NVRAM, use the following command:

```
clear log messages nvram
```

Example

The following command displays messages with a critical severity:

```
show log severity critical
```

The following command displays messages with warning, error, or critical severity:

```
show log severity warning
```

The following is sample output from a switch:

```
11/12/2004 00:38:10.30 <Warn:dm.Warn> MSM-A: Insufficient Power to power-on Slot-7
11/12/2004 00:38:08.77 <Warn:dm.Warn> MSM-A: Slot-7 being Powered OFF due to insuf
ficient power
11/12/2004 00:36:23.77 <Warn:dm.Warn> MSM-A: Slot-7 being Powered OFF due to insuf
ficient power
...
A total of 83 log messages were displayed.
```

The following command displays messages containing the string "slot 2":

```
show log match "slot 2"
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show log components

```
show log components {event component } {version}
```

Description

Displays the name, description and default severity for all components.

Syntax Description

<i>event component</i>	Specifies the component to display.
version	Specifies the version number of the component.

Default

N/A.

Usage Guidelines

This command displays the name, description, and default severity defined for the specified components or subcomponents.

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

Example

The following command displays the log components:

```
show log components
```

The following is sample output from this command:

Severity Component	Title	Threshold
AAA	Authentication, Authorization, Accounting	Info
RADIUS	Remote Authentication Dial In User Service	Error
TACACS	Terminal Access Controller Access Control Syst	Info
ACL	ACL	Info
CLEARFlow	CLEARFlow	Info
Policy	Policy actions	Info
bgp	Border Gateway Protocol	Info
damp	BGP Route Flap Dampening related debug message	Error
event	BGP FSM related events	Error
inUpdt	Incoming Update related debug msg	Warning
keepalive	BGP keepalive message	Warning
misc	Miscellaneous debug (Import, Aggregate, NextHop	Warning
msgs	Debug for BGP messages (OPEN, Update, Notifica	Warning
outUpdt	Transmit Update related debug	Warning
bootp	BOOTP, DHCP Component	Error
relay	BOOTP Relay trace component	Error
server	DHCP Server subcomponent	Info
cli	Command Line Interface	Info
shell	CLI configuration shell.	Error
subagent	CLI application subagent	Error
cm	Configuration Manager	Warning
file	CM file operation events	Warning
sys	CM system events	Warning
DM	Device Manager	Info
Card	Device Manager Card State Machine	Info
dosprot	dosprot	Info
ds	Directory Services	Error
EAPS	Ethernet Automatic Protection Switching	Info
SharedPort	EAPS SharedPort Domain	Info
EDP	Extreme DIScovery Protocol (EDP)	Error
ELRP		
Report	Extreme Loop Recognition Protocol	Warning
EPM	Extreme Process Manager	Info
KLM	Kernel Loadable Module Manager	Notice
Msg	Message Handler	Info
Upgrade	Upgrade Manager	Info
Version	Version Manager	Critical
ESRP	Extreme Standby Router Protocol	Error
Aware	Subsystem description	Info
InPdu	Subsystem description	Info
Nbr	Subsystem description	Info
OutPdu	Subsystem description	Info
State	ESRP State Transitions	Warning
System	Subsystem description	Warning
Track	Subsystem description	Warning
Vlan	Extreme Standby Router Protocol	Info
fdb	fdb module event	Error
HAL	Hardware Abstraction Layer	Error
Card	Card State Driver	Info
FDB	Forwarding Database Driver	Info
IPv4ACL	IPv4 Access Control List Driver	Info
IPv4Adj	IPv4 Adjacency Driver	Info

IPv4FIB	IPv4 FIB Driver	Info	
IPv4Mc	IPv4 Multicast Driver	Info	
Mirror	Mirroring Driver	Error	
Msg	Message Handler	Info	
Port	I/O Port Driver	Info	
SM	Switch Manager	Info	
Sys	System Driver	Info	
VLAN	VLAN Driver	Info	
IPMC	IP Multicast Main Module		Info
Snoop	IP Multicast Snooping Module	Error	
VLAN	IP Multicast VLAN Module	Error	
ISIS	Intermediate-to-Intermediate		Error
Export	Route Redistribution into ISIS	Error	
IFSM	ISIS Interface Finite State Machine (IFSM)	Warning	
IIH	ISIS Hello (IIH) PDU	Warning	
LSP	ISIS Link State PDU	Notice	
NFSM	ISIS Neighbor Finite State Machine (NFSM)	Warning	
PDU	ISIS General PDU	Warning	
Restart	ISIS Restart	Notice	
SPF	ISIS Shortest Path First (SPF)	Warning	
VLAN	ISIS VLAN-Related Events	Error	
Kern	Kernel messages		Error
LACP	Link Aggregation Control Protocol		Info
lldp	Link Layer Discovery Protocol (IEEE 802.1AB)		Warning
log	Log server messages		Warning
netTool	netTools framework		Error
dnsclient	Dns Client	Error	
dnsproxy	Dns Proxy	Error	
routeradv	IPv6 Router Advertisements	Warning	
sntp	Sntp client	Warning	
nl	Network Login		Info
dot1x	802.1X-based Network Login	Warning	
mac	MAC-based Network Login	Warning	
web	Web-based Network Login	Warning	
NM	Node Manager		Info
ospf	open shortest path first		Error
event	ospf events	Info	
hello	ospf hello	Error	
lsa	ospf link-state advertisement	Error	
neighbor	ospf neighbor	Error	
spf	ospf shortest path first	Error	
ospfv3	OSPFv3 related EMS messages		Warning
events	OSPF6 events related messages	Error	
lsa	LSA related messages	Warning	
nbr	OSPF6 neighbor related EMS messages	Warning	
pkt	OSPF6 Packet receive/transmit/processing relat	Warning	
route	OSPF6 route add/delete related messages	Warning	
spf	SPF computation related messages	Error	
pim	Pim Protocol Events		Warning
cache	PIM cache maintenance.	Warning	
debug	PIM debug messages	Notice	
hello	Hello messages	Warning	
mcdbg	multicast forwarding engine	Warning	
msg	Trace for pim control packtes	Notice	
nbr	Neighbor creation/deletion etc	Warning	
rpm	RP message exchange.	Warning	
pm	Policy Manager		Error
config	Policy file events	Info	
POE	Inline Power		Notice
rip	RIP routing		Error
cfg	rip configuration	Warning	
event	rip events	Warning	
inUpdt	rip - inbound route updates	Warning	
msg	rip - socket messages in and out	Warning	

outUpdt	rip - outbound route updates	Warning	
sys	rip - exos kernel interface	Warning	
ripng	RIPng Protocol Events		Warning
debug	RIPng debug messages	Notice	
external	RIPng external interface related messages	Warning	
message	RIPng control messages	Warning	
route	Hello messages	Warning	
rmon	RMON general info		Error
alarm	RMON alarm info	Error	
estat	RMON statistics info	Error	
event	RMON event info	Error	
history	RMON history	Error	
RtMgr	Route Manager		Info
VLAN	rtmgr vlan interface	Info	
sflow	Sflow Protocol Events		Warning
debug	SFLOW debug messages	Notice	
extended	SFLOW extended data collection	Notice	
msg	SFLOW process initializaion related message	Warning	
sample	SFLOW sample collection related messages	Warning	
statistics	SFLOW port statistics related message	Warning	
STP	Spanning-Tree Protocol		Error
InBPDU	STP In Bridge Protocol Data Unit	Warning	
OutBPDU	STP Out Bridge Protocol Data Unit	Warning	
System	STP System	Error	
System	XOS system related log messages		Info
telnetd	telnet server		Info
tftpd	tftp server		Info
thttpd	thttp server		Info
trace	Debug trace messages		Warning
vlan	Vlan mgr		Info
ack	vlan ack	Error	
dbg	Debug information	Info	
err	errors	Error	
mac	Virtual MAC Debugging	Info	
msgs	Messages	Info	
VRRP	Config/State messages		Warning
Advert	Subsystem description	Warning	
System	System/Library messages	Warning	
A total of 143 component(s) were displayed.			

The following command displays the version number of the VRRP component:

```
show log components vrrp version
```

The following is sample output from this command:

Component	Title	Version
VRRP	Config/State messages	2.4
Advert	Subsystem description	3.1
System	System/Library messages	3.2
A total of 3 component(s) were displayed.		

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show log configuration filter

```
show log configuration filter {filter_name}
```

Description

Displays the log configuration for the specified filter.

Syntax Description

<i>filter_name</i>	Specifies the filter to display.
--------------------	----------------------------------

Default

If no options are specified, the command displays the configuration for all filters.

Usage Guidelines

This command displays the configuration for filters.

Example

The following command displays the configuration for the filter, myFilter:

```
show log configuration filter myFilter
```

The following is sample output from this command:

```
Log Filter Name: myFilter
I/                                     Severity
E Comp.   Sub-comp.   Condition           CEWNISVD
- - - - -
I STP                                           -----
I aaa                                           -----
Include/Exclude: I - Include, E - Exclude
Component Unreg: * - Component/Subcomponent is not currently registered
Severity Values: C - Critical, E - Error, W - Warning, N - Notice, I - Info
* - Pre-assigned severities in effect for specified component
Debug Severity : S - Debug-Summary, V - Debug-Verbose, D - Debug-Data
+ - Debug Severities, but log debug-mode not enabled
If Match parameters present:
Parameter Flags: S - Source, D - Destination, (as applicable)
I - Ingress, E - Egress, B - BGP
Parameter Types: Port - Physical Port list, Slot - Physical Slot #
MAC - MAC address, IP - IP Address/netmask, Mask - Netmask
VID - Virtual LAN ID (tag), VLAN - Virtual LAN name
L4 - Layer-4 Port #, Num - Number, Str - String
Nbr - Neighbor, Rtr - Routerid, EAPS - EAPS Domain
Proc - Process Name
```

Strict Match : Y - every match parameter entered must be present in the event
 N - match parameters need not be present in the event

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show log configuration target

```
show log configuration target {console | memory-buffer | nvram |
    primary-node | backup-node | session | syslog {ipaddress {udp-port
    {udp_port }}| ipPort |ipaddress tls-port {tls_port}} {vr vr_name}
    {[local0...local7]}}
```

Description

Displays the log configuration for the specified target.

Syntax Description

console	Show the log configuration for the console display.
memory-buffer	Show the log configuration for volatile memory.
nvram	Show the log configuration for NVRAM.
primary-node	Specifies the primary node in a stack.
backup-node	Specifies the backup-node in a stack.
session	Show the log configuration for the current session (including console display).
syslog	Show the configuration for the specified syslog target.
<i>ipaddress</i>	Specifies the syslog IP address.
<i>ipPort</i>	Specifies the UDP port number for the syslog target.
tls_port	Specifies remote Syslog server Transport Layer Security (TLS) for connection type.
<i>tls_port</i>	TLS port number.
<i>vr_name</i>	Specifies the virtual router that can reach the server IP address. Note: User-created VRs are supported only on the platforms listed for this feature in the Switch Engine 32.3 Feature License Requirements document..
local0 ... local7	Specifies the local syslog facility.

Default

If no options are specified, the command displays the configuration for the current session and console display.

If a virtual router is not specified, VR-Mgmt is used.

Usage Guidelines

This command displays the log configuration for the specified target. The associated filter, severity, match expression, and format is displayed.

Example

The following command displays the log configuration:

```
show log configuration target
```

The following is sample output from this command:

```
Log Target      : memory-buffer
Enabled ?      : yes
Filter Name     : DefaultFilter
Match regex    : Any
Severity       : Debug-Data (through Critical)
Format        : MM/DD/YYYY HH:MM:SS.hh <Severity:Component.SubComponent.Condit
ion>
Buffer size    : 1000 messages
Log Target      : nvram
Enabled ?      : yes
Filter Name     : DefaultFilter
Match regex    : Any
Severity       : Warning (through Critical)
Format        : MM/DD/YYYY HH:MM:SS.hh <Severity:Component.SubComponent.Condit
ion>
Log Target      : console
Enabled ?      : no
Filter Name     : DefaultFilter
Match regex    : Any
Severity       : Info (through Critical)
Format        : MM/DD/YYYY HH:MM:SS.hh <Severity:Component.SubComponent.Condit
ion>
Log Target      : primary-msm
Enabled        : yes
Filter Name     : DefaultFilter
Match regex    : Any
Severity       : Warning (through Critical)
Log Target      : backup-msm
Enabled        : yes
Filter Name     : DefaultFilter
Match regex    : Any
Severity       : Warning (through Critical)
```

In the case that the alert is not configured:

```
# sh log configuration
Debug-Mode: Disabled

Log Target      : memory-buffer
Enabled        : yes
```

```

Filter Name : DefaultFilter
Match regex : Any
Severity : Debug-Data (through Critical)
Format : MM/DD/YYYY HH:MM:SS.hh <Severity:Component.SubComponent.Condition>
Buffer size : 1000 messages
Percent Full : 60%
Full Alert : None

Log Target : nvram
Enabled : yes
Filter Name : DefaultFilter
Match regex : Any
Severity : Warning (through Critical)
Format : MM/DD/YYYY HH:MM:SS.hh <Severity:Component.SubComponent.Condition>

Log Target : console
Enabled : no
Filter Name : DefaultFilter
Match regex : Any
Severity : Info (through Critical)
Format : MM/DD/YYYY HH:MM:SS.hh <Severity:Component.SubComponent.Condition>

```

In the case that the alert is configured and the percentage threshold is set to 90:

```

# sh log configuration
Debug-Mode: Disabled

Log Target : memory-buffer
Enabled : yes
Filter Name : DefaultFilter
Match regex : Any
Severity : Debug-Data (through Critical)
Format : MM/DD/YYYY HH:MM:SS.hh <Severity:Component.SubComponent.Condition>
Buffer size : 1000 messages
Percent Full : 60%
Full Alert : 90%

Log Target : nvram
Enabled : yes
Filter Name : DefaultFilter
Match regex : Any
Severity : Warning (through Critical)
Format : MM/DD/YYYY HH:MM:SS.hh <Severity:Component.SubComponent.Condition>

Log Target : console
Enabled : no
Filter Name : DefaultFilter
Match regex : Any
Severity : Info (through Critical)
Format : MM/DD/YYYY HH:MM:SS.hh <Severity:Component.SubComponent.Condition>

```

In the case that the buffer is currently wrapping and the percentage threshold is set to 90:

```

# show log configuration
Debug-Mode: Disabled

Log Target : memory-buffer
Enabled : yes
Filter Name : DefaultFilter
Match regex : Any
Severity : Debug-Data (through Critical)
Format : MM/DD/YYYY HH:MM:SS.hh <Severity:Component.SubComponent.Condition>
Buffer size : 1000 messages
Percent Full : 100% (wrapping)
Full Alert : 90%

```

```

Log Target      : nvram
  Enabled       : yes
  Filter Name   : DefaultFilter
  Match regex   : Any
  Severity      : Warning (through Critical)
  Format        : MM/DD/YYYY HH:MM:SS.hh <Severity:Component.SubComponent.Condition>

Log Target      : console
  Enabled       : no
  Filter Name   : DefaultFilter
  Match regex   : Any
  Severity      : Info (through Critical)
  Format        : MM/DD/YYYY HH:MM:SS.hh <Severity:Component.SubComponent.Condition>

```

Showing Syslog as log target:

```

# show log configuration target syslog

Log Target      : syslog; 10.68.6.3:6555 (vr VR-Mgmt), local0
  Enabled       : yes
  Filter Name   : DefaultFilter
  Match regex   : Any
  Severity      : Debug-Data (through Critical)
  Format        : PRI Mmm DD HH:MM:SS HOSTNAME TAG:
  Port Type     : TLS
  Recnct Cnt    : 2
  Recnct Msg    : CA Certificate not found. Use 'Download ssl certificate trusted-ca'
                  command to download a CA certificate.

Log Target      : syslog; 10.68.6.3:6555 (vr VR-Mgmt), local1
  Enabled       : no
  Filter Name   : DefaultFilter
  Match regex   : Any
  Severity      : Debug-Data (through Critical)
  Format        : PRI Mmm DD HH:MM:SS HOSTNAME TAG:
  Port Type     : TLS
  Recnct Cnt    : 0
  Recnct Msg    : No Error

Log Target      : syslog; 10.68.6.3:6519 (vr VR-Mgmt), local0
  Enabled       : yes
  Filter Name   : DefaultFilter
  Match regex   : Any
  Severity      : Debug-Data (through Critical)
  Format        : PRI Mmm DD HH:MM:SS HOSTNAME TAG:
  Port Type     : UDP

```

History

This command was first available in ExtremeXOS 10.1.

The **ipPort** parameter was first available in ExtremeXOS 11.0.

The **local0 ... local17** keywords were made optional in ExtremeXOS 11.4.

The **udp-port** parameter was added in ExtremeXOS 21.1.

Transport Layer Security (TLS) option added in ExtremeXOS 22.1.

Connection port type (TLS or UDP) and connection messages added in ExtremeXOS 22.1

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show log configuration target upm

```
show log configuration target upm {upm_profile_name}
```

Description

Displays a UPM target profile configuration.

Syntax Description

<i>upm_profile_name</i>	Specifies the name of the UPM target profile you want to view.
-------------------------	--

Default

N/A.

Usage Guidelines

None.

Example

The following example displays the configuration for the UPM log target named testprofile1:

```
show log configuration target upm testprofile1
```

In the case that the alert is not configured:

```
X670V-48t.1 # sh log configuration
Debug-Mode: Disabled

Log Target      : memory-buffer
Enabled        : yes
Filter Name     : DefaultFilter
Match regex    : Any
Severity       : Debug-Data (through Critical)
Format         : MM/DD/YYYY HH:MM:SS.hh <Severity:Component.SubComponent.Condition>
Buffer size    : 1000 messages
Percent Full   : 60%
Full Alert     : None

Log Target      : nvram
Enabled        : yes
Filter Name     : DefaultFilter
Match regex    : Any
Severity       : Warning (through Critical)
```

```

Format      : MM/DD/YYYY HH:MM:SS.hh <Severity:Component.SubComponent.Condition>

Log Target   : console
Enabled      : no
Filter Name  : DefaultFilter
Match regex  : Any
Severity     : Info (through Critical)
Format      : MM/DD/YYYY HH:MM:SS.hh <Severity:Component.SubComponent.Condition>

```

In the case that the alert is configured and the percentage threshold is set to 90:

```

X670V-48t.1 # sh log configuration
Debug-Mode: Disabled

Log Target   : memory-buffer
Enabled      : yes
Filter Name  : DefaultFilter
Match regex  : Any
Severity     : Debug-Data (through Critical)
Format      : MM/DD/YYYY HH:MM:SS.hh <Severity:Component.SubComponent.Condition>
Buffer size  : 1000 messages
Percent Full : 60%
Full Alert   : 90%

Log Target   : nvram
Enabled      : yes
Filter Name  : DefaultFilter
Match regex  : Any
Severity     : Warning (through Critical)
Format      : MM/DD/YYYY HH:MM:SS.hh <Severity:Component.SubComponent.Condition>

Log Target   : console
Enabled      : no
Filter Name  : DefaultFilter
Match regex  : Any
Severity     : Info (through Critical)
Format      : MM/DD/YYYY HH:MM:SS.hh <Severity:Component.SubComponent.Condition>

```

In the case that the buffer is currently wrapping and the percentage threshold is set to 90:

```

X670V-48t.1 # sh log configuration
Debug-Mode: Disabled

Log Target   : memory-buffer
Enabled      : yes
Filter Name  : DefaultFilter
Match regex  : Any
Severity     : Debug-Data (through Critical)
Format      : MM/DD/YYYY HH:MM:SS.hh <Severity:Component.SubComponent.Condition>
Buffer size  : 1000 messages
Percent Full : 100% (wrapping)
Full Alert   : 90%

Log Target   : nvram
Enabled      : yes
Filter Name  : DefaultFilter
Match regex  : Any
Severity     : Warning (through Critical)
Format      : MM/DD/YYYY HH:MM:SS.hh <Severity:Component.SubComponent.Condition>

Log Target   : console
Enabled      : no
Filter Name  : DefaultFilter

```

```
Match regex : Any
Severity    : Info (through Critical)
Format      : MM/DD/YYYY HH:MM:SS.hh <Severity:Component.SubComponent.Condition>
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the Universal Port feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

show log configuration target xml-notification

```
show log configuration target xml-notification {xml_target_name}
```

Description

Displays XML target information.

Syntax Description

<code>xml_target_name</code>	Specifies the configured xml notification target.
------------------------------	---

Default

N/A.

Usage Guidelines

Use this command to display XML target information.

Example

The following command displays XML target information for all targets:

```
show log configuration target xml-notification
```

Following is sample output from the command:

```
Log Target      : xml-notification (sqa)
Enabled        : yes
Filter Name     : DefaultFilter
Match regex    : Any
Severity       : Info (through Critical)
Log Target      : xml-notification (epi)
Enabled        : yes
Filter Name     : xmlc_filter_epi
```

```
Match regex : Any
Severity    : Info (through Critical)
```

In the case that the alert is not configured:

```
X670V-48t.1 # sh log configuration
Debug-Mode: Disabled

Log Target      : memory-buffer
  Enabled       : yes
  Filter Name   : DefaultFilter
  Match regex   : Any
  Severity      : Debug-Data (through Critical)
  Format        : MM/DD/YYYY HH:MM:SS.hh <Severity:Component.SubComponent.Condition>
  Buffer size    : 1000 messages
  Percent Full  : 60%
  Full Alert    : None

Log Target      : nvram
  Enabled       : yes
  Filter Name   : DefaultFilter
  Match regex   : Any
  Severity      : Warning (through Critical)
  Format        : MM/DD/YYYY HH:MM:SS.hh <Severity:Component.SubComponent.Condition>

Log Target      : console
  Enabled       : no
  Filter Name   : DefaultFilter
  Match regex   : Any
  Severity      : Info (through Critical)
  Format        : MM/DD/YYYY HH:MM:SS.hh <Severity:Component.SubComponent.Condition>
```

In the case that the alert is configured and the percentage threshold is set to 90:

```
X670V-48t.1 # sh log configuration
Debug-Mode: Disabled

Log Target      : memory-buffer
  Enabled       : yes
  Filter Name   : DefaultFilter
  Match regex   : Any
  Severity      : Debug-Data (through Critical)
  Format        : MM/DD/YYYY HH:MM:SS.hh <Severity:Component.SubComponent.Condition>
  Buffer size    : 1000 messages
  Percent Full  : 60%
  Full Alert    : 90%

Log Target      : nvram
  Enabled       : yes
  Filter Name   : DefaultFilter
  Match regex   : Any
  Severity      : Warning (through Critical)
  Format        : MM/DD/YYYY HH:MM:SS.hh <Severity:Component.SubComponent.Condition>

Log Target      : console
  Enabled       : no
  Filter Name   : DefaultFilter
  Match regex   : Any
  Severity      : Info (through Critical)
  Format        : MM/DD/YYYY HH:MM:SS.hh <Severity:Component.SubComponent.Condition>
```

In the case that the buffer is currently wrapping and the percentage threshold is set to 90:

```
X670V-48t.1 # sh log configuration
Debug-Mode: Disabled

Log Target      : memory-buffer
Enabled         : yes
Filter Name     : DefaultFilter
Match regex     : Any
Severity        : Debug-Data (through Critical)
Format          : MM/DD/YYYY HH:MM:SS.hh <Severity:Component.SubComponent.Condition>
Buffer size     : 1000 messages
Percent Full    : 100% (wrapping)
Full Alert      : 90%

Log Target      : nvram
Enabled         : yes
Filter Name     : DefaultFilter
Match regex     : Any
Severity        : Warning (through Critical)
Format          : MM/DD/YYYY HH:MM:SS.hh <Severity:Component.SubComponent.Condition>

Log Target      : console
Enabled         : no
Filter Name     : DefaultFilter
Match regex     : Any
Severity        : Info (through Critical)
Format          : MM/DD/YYYY HH:MM:SS.hh <Severity:Component.SubComponent.Condition>
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show log configuration

```
show log configuration
```

Description

Displays the log configuration for switch log settings, and for certain targets.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command displays the log configuration for all targets. The state of the target, enabled or disabled, appears. For the enabled targets, the associated filter, severity, match expression, and format appears. The debug mode state of the switch is also appears.

Example

The following command displays the configuration of all the log targets and all existing filters:

```
# show log configuration
Debug-Mode: Disabled
Show Message Privilege: User (non-administrative)
Enabled TLS Ciphers: aes128-sha256, dhe-rsa-aes256-sha256
Disabled TLS Ciphers: aes128-sha, aes256-sha256, dhe-rsa-aes128-sha256
Syslog TLS TCP User Timeout: Default
Syslog TLS OCSP: On

Log Target      : memory-buffer
  Enabled       : yes
  Filter Name   : DefaultFilter
  Match regex   : Any
  Severity      : Debug-Data (through Critical)
  Format        : MM/DD/YYYY HH:MM:SS.hh <Severity:Component.SubComponent.Condition>
  Buffer size    : 1000 messages
  Percent Full  : 100% (wrapping)
  Full Alert    : None

Log Target      : nvram
  Enabled       : yes
  Filter Name   : DefaultFilter
  Match regex   : Any
  Severity      : Warning (through Critical)
  Format        : MM/DD/YYYY HH:MM:SS.hh <Severity:Component.SubComponent.Condition>

Log Target      : console
  Enabled       : no
  Filter Name   : DefaultFilter
  Match regex   : Any
  Severity      : Info (through Critical)
  Format        : MM/DD/YYYY HH:MM:SS.hh <Severity:Component.SubComponent.Condition>

Log Filter Name: DefaultFilter
I/              Severity
E Component     SubComponent Condition      CEWNISVD
- - - - -
I All           *****

Include/Exclude: I - Include, E - Exclude
Component Unreg: * - Component/SubComponent is not currently registered
Severity Values: C - Critical, E - Error, W - Warning, N - Notice, I - Info
                  * - Pre-assigned severities in effect for specified component
Debug Severity : S - Debug-Summary, V - Debug-Verbose, D - Debug-Data
                  + - Debug Severities, but log debug-mode not enabled
If Match parameters present:
Parameter Flags: S - Source, D - Destination, (as applicable)
                  I - Ingress, E - Egress, B - BGP
Parameter Types: Port - Physical Port list, Slot - Physical Slot #
                  MAC - MAC address, IP - IP Address/netmask, Mask - Netmask
                  VID - Virtual LAN ID (tag), VLAN - Virtual LAN name
                  VR - Virtual Router Name, VRID - VR Identifier
                  VRF - Virtual Routing and Forwarding Name
```

```

L4 - Layer-4 Port #, Num - Number, Str - String
Nbr - Neighbor, Rtr - Routerid, EAPS - EAPS Domain
Proc - Process Name
Strict Match : Y - every match parameter entered must be present in the event
               N - match parameters need not be present in the event

```

The following command displays Syslog TLS OCSP attributes (lines 8-11):

```

# show log configuration
Debug-Mode: Disabled
Show Message Privilege: User (non-administrative)
Enabled TLS Ciphers:  dhe-rsa-aes128-sha256, aes256-sha256, dhe-rsa-aes256-sha256, aes128-
sha, aes128-sha256
Disabled TLS Ciphers: None
Syslog TLS TCP User Timeout: Default
Syslog TLS OCSP: On
Syslog TLS OCSP Attributes:
  Nonce                : On
  Signer ocsp-nocheck : On
  Override Server URL : http://syslogocsp:2022

Log Target      : memory-buffer
  Enabled       : yes
  Filter Name   : DefaultFilter
  Match regex   : Any
  Severity      : Debug-Data (through Critical)
  Format        : MM/DD/YYYY HH:MM:SS.hh <Severity:Component.SubComponent.Condition>
  Buffer size    : 1000 messages
  Percent Full  : 100% (wrapping)
  Full Alert    : None

Log Target      : nvram
  Enabled       : yes
  Filter Name   : DefaultFilter
  Match regex   : Any
  Severity      : Warning (through Critical)
  Format        : MM/DD/YYYY HH:MM:SS.hh <Severity:Component.SubComponent.Condition>

Log Target      : console
  Enabled       : no
  Filter Name   : DefaultFilter
  Match regex   : Any
  Severity      : Info (through Critical)
  Format        : MM/DD/YYYY HH:MM:SS.hh <Severity:Component.SubComponent.Condition>

Log Filter Name: DefaultFilter
I/
E Component    SubComponent Condition          Severity
- - - - -
I All
*****

Include/Exclude: I - Include, E - Exclude
Component Unreg: * - Component/SubComponent is not currently registered
Severity Values: C - Critical, E - Error, W - Warning, N - Notice, I - Info
                  * - Pre-assigned severities in effect for specified component
Debug Severity : S - Debug-Summary, V - Debug-Verbose, D - Debug-Data
                  + - Debug Severities, but log debug-mode not enabled
If Match parameters present:
Parameter Flags: S - Source, D - Destination, (as applicable)
                  I - Ingress, E - Egress, B - BGP
Parameter Types: Port - Physical Port list, Slot - Physical Slot #
                  MAC - MAC address, IP - IP Address/netmask, Mask - Netmask
                  VID - Virtual LAN ID (tag), VLAN - Virtual LAN name

```

```

VR    - Virtual Router Name, VRID - VR Identifier
VRF   - Virtual Routing and Forwarding Name
L4    - Layer-4 Port #, Num - Number, Str - String
Nbr   - Neighbor, Rtr - Routerid, EAPS - EAPS Domain
Proc  - Process Name
Strict Match : Y - every match parameter entered must be present in the event
              N - match parameters need not be present in the event

```

History

This command was first available in ExtremeXOS 10.1.

Syslog reference identifier information was added in ExtremeXOS 22.3.

Which Syslog TLS session ciphers are enabled/disabled was added in ExtremeXOS 22.4.

Information about the Syslog TLS TCP user timeout value was added in ExtremeXOS 22.5.

OCSP check status information was added in ExtremeXOS 30.7.

Syslog TLS OCSP attributes were added in ExtremeXOS 32.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show log counters

```

show log counters {event condition | [all | event component]} {include |
notified | occurred} {severity severity {only}}

```

Description

Displays the incident counters for events.

Syntax Description

<i>event condition</i>	Specifies the event condition to display.
all	Specifies that all events are to be displayed.
<i>event component</i>	Specifies that all the events associated with a particular component or subcomponent should be displayed.
include	Specifies if one or more targets should be included in this event.
notified	Specifies the number of times this event has occurred.
occurred	Specifies the number of times this event has occurred since the last clear or reboot.
<i>severity</i>	Specifies the minimum severity level of events to display (if the keyword only is omitted).
only	Specifies that only events of the specified severity level are to be displayed.

Default

If severity is not specified, then events of all severity are displayed.

Usage Guidelines

This command displays the incident counters for each event specified. Two incident counters are displayed. One counter displays the number of times an event has occurred, and the other displays the number of times that notification for the event was made to the system (an incident record was injected into the system for further processing). Both incident counters reflect totals accumulated since reboot or since the counters were cleared using the `clear log counters` or `clear counters` command, regardless of whether it was filtered or not.

The keywords include, notified, and occurred only display events with non-zero counter values for the corresponding counter.

This command also displays a reference count (the column titled Rf in the output). The reference count is the number of enabled targets receiving notifications of this event.

See the command `show log` for more information about severity levels.

To get a listing of the event conditions in the system, use the following command:

```
show log events
```

To get a listing of the components present in the system, use the following command:

```
show log components
```

Example

The following example displays the event counters for event conditions of severity debug-summary or greater in the component STP:

```
# show log counters stp included severity debug-summary
```

Component	SubComponent	Condition	Severity	Occurred	I	Ntfd	Last
Notified Time							

STP		AllocClntMsgFail	Warning	0	Y	0	
STP		BindAutobindFail	Error	0	Y	0	
STP		ChangeTag	Warning	0	Y	0	
STP		CreatDmnFail	Warning	0	Y	0	
STP		DomainDisable	Warning	0	Y	0	
STP		DomainEnable	Warning	1	Y	1	01/05/2001
23:43:39.03							
STP		DsblPortBrdgDtect	Warning	0	Y	0	
STP		DsblPortLoopDtect	Warning	0	Y	0	
STP		EnblPortTimerExp	Warning	0	Y	0	
STP		ESRPCConflict	Error	0	Y	0	
STP		GetAvailDomainFail	Error	0	Y	0	
STP		GetAvailPortFail	Error	0	Y	0	
STP		LicenseError	Warning	0	Y	0	
STP		NULLBridgeId	Error	0	Y	0	
STP		PhysPortAssocInv	Warning	0	Y	0	
STP		PhysPortInv	Error	0	Y	0	
STP		PortAlrdyBndDmn	Warning	0	Y	0	
STP		PortInuse	Error	0	Y	0	

STP	PortModeNotExpct	Error	0	Y	0
STP	SendClntMsgFail	Warning	0	Y	0
STP	UnsupPhysInterface	Critical	0	Y	0

The following example displays the event counters for the event condition PDUDrop in the component STP.InBPDU:

```
# show log counters "STP.InBPDU.Drop"
Component SubComponent Condition Severity Occurred I Ntfd Last
Notified Time
-----
STP InBPDU Drop Error 0 Y 0

Occurred : # of times this event has occurred since last clear or reboot
Flags : (*) Not all applications responded in time with their counters
I(ncluded): Set to Y(es) if one or more targets filter includes this event
Notified : # of times this event has occurred when 'Included' was Y(es)
```

The following example displays the event counters for the AAA component:

```
# show log counters aaa
Component SubComponent Condition Severity Occurred I Ntfd Last
Notified Time
-----
AAA accountLockedOut Warning 0 Y 0
AAA accountMod Info 0 Y 0
AAA authFail Warning 0 Y 0
AAA authPass Info 2 Y 2 02/27/2018
16:27:10.80
AAA changePass Info 0 Y 0
AAA ChgAccntPrvlgFail Notice 0 Y 0
AAA ChgAccntPrvlgOK Notice 0 Y 0
AAA ClearSessAccntDel Notice 0 Y 0
AAA ClearSessAccntPrvlgChg Notice 0 Y 0
AAA createAccount Info 0 Y 0
AAA CreateEncryptAccnt Info 0 Y 0
AAA DecryptMaxStrSizExcd Warning 0 Y 0
AAA deleteAccount Info 0 Y 0
AAA EnblFIPSMoDeFail Error 0 Y 0
AAA EnblFIPSMoDeOK Notice 0 Y 0
AAA InitRAccSvVr Warning 0 Y 0
AAA InitRAutSvVr Warning 0 Y 0
AAA InitVlanLibFail Error 0 Y 0
AAA ipv6ForceLocal Debug-Summary 0 N 0
AAA localAuthen Debug-Summary 2 N 0
AAA logout Info 2 Y 2 02/27/2018
16:06:58.31
```

History

This command was first available in ExtremeXOS 10.1.

Last notified information was added in ExtremeXOS 22.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show log events

```
show log events [event condition | [all | event component] {severity
    severity {only}}] {details}
```

Description

Displays information about the individual events (conditions) that can be logged.

Syntax Description

<i>event condition</i>	Specifies the event condition to display.
all	Specifies that all events are to be displayed.
<i>event component</i>	Specifies that all the events associated with a particular component should be displayed.
<i>severity</i>	Specifies the minimum severity level of events to display (if the keyword only is omitted).
only	Specifies that only events of the specified severity level are to be displayed.
details	Specifies that detailed information, including the message format and parameter types, be displayed.

Default

If severity is not specified, then events of all severity are displayed. If detail is not specified, then summary only information is displayed.

Usage Guidelines

This command displays the mnemonic, message format, severity, and parameter types defined for each condition in the event set specified.

See the command [show log](#) for more information about severity levels.

When the detail option is specified, the message format is displayed for the event conditions specified. The message format parameters are replaced by the value of the parameters when the message is generated.

To get a listing of the components present in the system, use the following command:

```
show log components
```

Example

The following command displays the event conditions of severity debug-summary or greater in the component STP.InBPDU:

```
show log events stp.inbpdu severity debug-summary
```

The following is sample output from this command:

Comp	SubComp	Condition	Severity	Parameters
STP	InBPDU	Drop	Error	2 total
STP	InBPDU	Ign	Debug-Summary	2 total
STP	InBPDU	Mismatch	Warning	2 total

The following command displays the details of the event condition PDUTrace in the component STP.InBPDU:

```
show log events stp.inbpdu.pdutrace details
```

The following is sample output from this command:

Comp	SubComp	Condition	Severity	Parameters
STP	InBPDU	Trace	Debug-Verbose	2 total
0 - string				
1 - string (printf)				
Port=%0%: %1%				

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show mac-lockdown-timeout fdb ports

```
show mac-lockdown-timeout fdb ports [all | port_list]
```

Description

Displays the MAC entries that are learned on the specified port or group of ports or for all ports on the switch along with the aging time of each port.

Syntax Description

all	Specifies all ports
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

N/A.

Usage Guidelines

If a port is down, the command displays all of the MAC entries that are maintained locally in the software.

The MAC entries learned on the specified port are displayed only if the MAC lock down timeout feature is enabled on the port. If you specify a port on which this feature is disabled, the MAC entries learned on that port are not displayed.

The switch displays the following information:

- Mac—The MAC address that defines the entry.
- Vlan—The *VLAN* name and ID for the entry.
- Age—The age of the entry, in seconds.
- Flags—Flags that define the type of entry:
 - B—Egress Blackhole.
 - b—Ingress Blackhole.
 - F—Entry in the hardware *FDB*.
 - L—Entry in the software.
- Port—The port on which the MAC address has been learned.

Example

The following command displays information about the MAC address lock down timeout settings for ports 2:3 and 2:4:

```
show mac-lockdown-timeout fdb ports 2:3, 2:4
```

The following is sample output from this command:

Mac	Vlan	Age	Flags	Port

00:00:01:02:03:04v1 (4094)	0010F2:3			
00:00:01:00:00:02v1 (4094)	0030FB b2:3			
00:00:0A:02:03:04v2 (4093)	0050L2:4			
00:00:0B:02:03:04v2 (4093)	0090F2:4			
Flags : (F) Entry as in h/w FDB, (L) Entry in s/w and not in h/w				
(B) Egress Blackhole, (b) Ingress Blackhole				
Total: 4 Entries in FDB: 3Entries in s/w: 1				

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show mac-lockdown-timeout ports

```
show mac-lockdown-timeout ports [all | port_list]
```

Description

Displays information about the MAC address lock down timeout feature for the specified port or group of ports or for all ports on the switch.

Syntax Description

all	Specifies all ports
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

N/A.

Usage Guidelines

The switch displays the following MAC address timeout information:

- Port—Indicates the port number that you specified in the command.
- MAC Lockdown Timeout—Specifies the enabled/disabled state of the MAC address lock down timeout feature.
- Timeout (in seconds)—Specifies the timeout value for the specified ports. By default, the timeout value is 15 seconds. Even if MAC address lock down is disabled, the default timeout value is displayed.

Example

The following command displays information about the MAC address lock down timeout settings for ports 2:3, 2:4, and 2:6:

```
show mac-lockdown-timeout ports 2:3, 2:4, 2:6
```

The following is sample output from this command:

```
Ports    MAC Lockdown Timeout    Timeout (in seconds)
=====
2:3 Enabled300
2:4 Enabled 300
2:6Disabled                15
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show mac-locking stations

```
show mac-locking {first-arrival | static} {ports port_list}
```

Description

Displays MAC locking information about end stations connected to the switch.

Syntax Description

first-arrival	Displays only first-arrival MAC locking information.
static	Displays only static MAC locking information.
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

N/A

Usage Guidelines

If you do not specify any parameters, MAC locking information for end stations on all ports (both static and first-arrival) is displayed.

Example

The following example displays MAC locking information for end stations on all ports.

```
show mac-locking stations

Port    MAC Address           Status   State           Aging
-----  -
1:2     00:00:11:22:33:00    inactive static          false
2:3     00:00:11:22:33:99    active  static          false
3:4     00:bb:00:00:00:00    active  first-arrival   false
4:5     00:e0:2b:00:00:01    active  first-arrival   true
Total for all ports: 4 Static: 2 First-Arrival: 2
```

History

This command was first available in ExtremeXOS 15.7.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show mac-locking

```
show mac-locking {ports port_list}
```

Description

Displays the status of MAC locking on one or more ports

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
------------------	---

Default

N/A

Usage Guidelines

If you do not specify a port, MAC locking status is displayed for all ports.

In MLAG, the mac-locking entries shown in this command's output are only natively learned FDB entries on the switch.

Example

The following example displays MAC locking status for all ports.

```
show mac-locking

MAC locking is globally disabled.
```

Port	MAC Lock Stat	Trap Thr	Log Viol	FA Thr	Viol	FA Aging	Limit Action Cfg	Link Down Stat	Action	Max Stc	Max FA	Last Violating MAC Address
1	dis	off	off	off	off	dis	ena ena	clear		64	600	00:00:00:00:00:00
2	dis	off	off	off	off	dis	ena ena	clear		64	600	00:00:00:00:00:00
3	dis	off	off	off	off	dis	ena ena	clear		64	600	00:00:00:00:00:00
4	dis	off	off	off	off	dis	ena ena	clear		64	600	00:00:00:00:00:00
5	dis	off	off	off	off	dis	ena ena	clear		64	600	00:00:00:00:00:00
6	dis	off	off	off	off	dis	ena ena	clear		64	600	00:00:00:00:00:00
7	dis	off	off	off	off	dis	ena ena	clear		64	600	00:00:00:00:00:00
8	dis	off	off	off	off	dis	ena ena	clear		64	600	00:00:00:00:00:00
9	dis	off	off	off	off	dis	ena ena	clear		64	600	00:00:00:00:00:00
10	dis	off	off	off	off	dis	ena ena	clear		64	600	00:00:00:00:00:00
11	dis	off	off	off	off	dis	ena ena	clear		64	600	00:00:00:00:00:00
12	dis	off	off	off	off	dis	ena ena	clear		64	600	00:00:00:00:00:00
13	dis	off	off	off	off	dis	ena ena	clear		64	600	00:00:00:00:00:00
14	dis	off	off	off	off	dis	ena ena	clear		64	600	00:00:00:00:00:00
15	dis	off	off	off	off	dis	ena ena	clear		64	600	00:00:00:00:00:00
16	dis	off	off	off	off	dis	ena ena	clear		64	600	00:00:00:00:00:00
17	dis	off	off	off	off	dis	ena ena	clear		64	600	00:00:00:00:00:00
18	dis	off	off	off	off	dis	ena ena	clear		64	600	00:00:00:00:00:00
19	dis	off	off	off	off	dis	ena ena	clear		64	600	00:00:00:00:00:00
20	dis	off	off	off	off	dis	ena ena	clear		64	600	00:00:00:00:00:00

```

21    dis  off|off  off|off  dis  ena|ena  clear  64 600  00:00:00:00:00:00
22    dis  off|off  off|off  dis  ena|ena  clear  64 600  00:00:00:00:00:00
23    dis  off|off  off|off  dis  ena|ena  clear  64 600  00:00:00:00:00:00
24    dis  off|off  off|off  dis  ena|ena  clear  64 600  00:00:00:00:00:00
25    dis  off|off  off|off  dis  ena|ena  clear  64 600  00:00:00:00:00:00
26    dis  off|off  off|off  dis  ena|ena  clear  64 600  00:00:00:00:00:00
27    dis  off|off  off|off  dis  ena|ena  clear  64 600  00:00:00:00:00:00
28    dis  off|off  off|off  dis  ena|ena  clear  64 600  00:00:00:00:00:00
29    dis  off|off  off|off  dis  ena|ena  clear  64 600  00:00:00:00:00:00
30    dis  off|off  off|off  dis  ena|ena  clear  64 600  00:00:00:00:00:00
31    dis  off|off  off|off  dis  ena|ena  clear  64 600  00:00:00:00:00:00
32    dis  off|off  off|off  dis  ena|ena  clear  64 600  00:00:00:00:00:00
33    dis  off|off  off|off  dis  ena|ena  clear  64 600  00:00:00:00:00:00
34    dis  off|off  off|off  dis  ena|ena  clear  64 600  00:00:00:00:00:00

```

Legend:

Stat	- Status	Thr Viol	- Threshold Violation
Max Stc	- Max Static Count	Max FA	- Max First-Arrival Count
dis	- Disabled	ena	- Enabled
retain	- Retain MACs	clear	- Clear MACs
Limit Action Cfg	- If port should be disabled when learnt limit is exceeded		
dis	- Port to be disabled when learn limit is exceeded		
ena	- Port to remain enabled when learn limit is exceeded		
Limit Action Stat	- Port status on exceeding learn limit		

History

This command was first available in ExtremeXOS 15.7.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show macsec

```
show macsec
```

Description

Displays a system-wide view of MAC Security (MACsec).

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command allows you to quickly see which ports support MACsec, which are configured, and which are secure.

- MACsec Capable without External Adapter—Ports that inherently support MACsec
- HW-Mode MACsec—Ports configured for MACsec versus for half-duplex (only applicable on half-duplex/MACsec ports).
- MACsec Capable with External Adapter—Ports that support MACsec-capable adapters.
- LRM/MACsec Adapter Present—Ports with a LRM/MACsec adapter plugged in.
- Valid MACsec License—Ports with a valid MACsec license installed.
- MACsec Capable, Present, and Licensed—Ports that support MACsec, external adapter is present (if applicable), and are licensed for MACsec.
- MACsec Configured—Ports that have been assigned to a connectivity association (CA) that in turn has been configured with a pre-shared-key (PSK).
- MKA Active—Ports that have MACsec configured and are actively participating in MKA (transmitting MKPDUs).
- Connect Status:
 - Pending—no connectivity (MKA not successful; no connectivity).
 - Authenticated—unsecure connectivity (peer authenticated; packets not encrypted).



Note

Extreme Network switches always attempt to connect securely. However, if the peer is a third-party device and the peer is elected key server and the peer chooses to connect without MACsec protection, the port's connect status becomes "authenticated" instead of "secure". In authenticated mode, MKA continues to authenticate the remote peer, but MACsec protection is not enabled and all traffic transmits in the clear.

- Secure—secure connectivity (peer authenticated, and packets encrypted).

For ports with shared media (one copper and one fiber), normally fiber is the preferred medium; however, for proper detection/operation, the fiber port must be the preferred medium. For example, if link is detected on the copper port it becomes the preferred medium. As such it is removed from the MACsec-capable port list. The copper ports of the shared media ports are not MACsec-capable. Only the fiber side with an LRM/MACsec adapter installed is MACsec-capable.

Example

The following example shows system-wide view of MACsec:

```
# show macsec
MACsec Capable Without External Adapter: 1:25-48,2:25-48
HW-Mode MACsec: 1:25-48,2:25-48
MACsec Capable with External Adapter: 1:49-54,2:49-54
LRM/MACsec Adapter Present: 2:49-50
Valid MACsec License: 1:25-54,2:25-54
MACsec Capable, Present and Licensed: 1:25-48,2:25-50
MACsec Configured: 1:37,1:48,2:25,2:29,2:32,2:49
MKA Active: 1:37,2:49 (Transmitting MKPDUs)
Connect Status
  Pending: 1:48,2:25,2:29,2:32 (No connectivity)
  Secure: 1:37,2:49 (Secured connectivity: MKA with
MACsec)
```

History

This command was first available in ExtremeXOS 30.1.

Platform Availability

This command is available on the following platforms.



Note

The MACsec feature requires the installation of the MAC Security feature pack license.

Platform	Ports
ExtremeSwitching 5320	All ports of all models except stacking ports.
ExtremeSwitching 5420	All ports of all models except stacking ports.
ExtremeSwitching 5520	All ports, except 5520-VIM-4X and 5520-24X 10G ports
ExtremeSwitching 5720	All ports of all models except stacking ports.

show macsec connectivity-association

```
show macsec { connectivity-association { ca_name }
```

Description

Displays a global summary of MAC Security (MACsec) capabilities and status for all or a specified connectivity-association (CA).

Syntax Description

connectivity-association	Secure connectivity provided between MACsec stations.
<i>ca_name</i>	Selects the Connectivity Association (CA) name.

Default

N/A.

Example

The following example shows general information about CAs:

```
# show macsec connectivity-association
MACsec
CA Name                               Ports                               CAK Bit
-----                               -
CentOS-50                             50                                256
foo128                                None                               128
foo256                                None                               256
MyLittleCa                             33                                128
MyBigCa                                None                               256
                                     CAK Name (CKN)
-----
```

The following example shows information about CA "testca":

```
# show macsec connectivity-association testca

MACsec Connectivity Association: ca25
  Pre-shared-key
    CKN: blue
    CAK: 256-bit
  Ports: 25
```

History

This command was first available in ExtremeXOS 30.1.

Information for 256 cipher suite support was added in ExtremeXOS 30.2.

Platform Availability

This command is available on the following platforms.



Note

The MACsec feature requires the installation of the MAC Security feature pack license.

Platform	Ports
ExtremeSwitching 5320	All ports of all models except stacking ports.
ExtremeSwitching 5420	All ports of all models except stacking ports.
ExtremeSwitching 5520	All ports, except 5520-VIM-4X and 5520-24X 10G ports
ExtremeSwitching 5720	All ports of all models except stacking ports.

show macsec encryption-engine monitor

```
show macsec encryption-engine-monitor
```

Description

Displays a global summary of the MAC Security (MACsec) Encryption Engine Monitor.

Syntax Description

encryption-engine-monitor	Displays the MACsec Encryption Engine Monitor.
----------------------------------	--

Default

N/A.

Example

The following example shows details of the Encryption Engine Monitor:

```
# show macsec encryption-engine-monitor
Encryption Engine Monitor
-----
Slots          Engine-1          Engine-2
          Packet Drops      Packet Drops
=====
Slot-1          74565          480989
Slot-2           341           0
Slot-3          n/a          n/a
Slot-4          2374          879374
Slot-5          n/a          n/a
Slot-6           0           0
Slot-7          23289          0
Slot-8           0          98740
```

History

This command was first available in ExtremeXOS 31.4.

Platform Availability

This command is available on the following platforms.



Note

The MACsec feature requires the installation of the MAC Security feature pack license.

Platform	Ports
ExtremeSwitching 5320	All ports of all models except stacking ports.
ExtremeSwitching 5420	All ports of all models except stacking ports.
ExtremeSwitching 5520	All ports, except 5520-VIM-4X and 5520-24X 10G ports
ExtremeSwitching 5720	All ports of all models except stacking ports.

show macsec ports

```
show macsec ports port-list usage
```

Description

Displays per-port MKA and MAC Security (MACsec) data in tabular format.

Syntax Description

ports	Specifies ports to show information on.
<i>port_list</i>	Lists which ports to view MACsec information on.
usage	Specifies to display per-port MACsec usage information.

Default

N/A.

Usage Guidelines

This command displays a table containing both control-layer (MKA) status and data-layer (MACsec) statistics:

- **Port**—Underlying physical port's name. Only MACsec capable ports appear.
- **MKA**—Shows the message number (MN) contained in the MKPDUs sent by the port ("Local MN"), as well as the MN's in the MKPDUs being received ("Peer MN"). During normal operation, each MN should increment by 1 once every 2 seconds (MKA Hello Time).
- **Peer Status**—Indicates whether or not the peer is potential or live. Per *IEEE802.1X-2010's Clause 9.4.3 Determining Liveness*, a peer is considered "live" when it transmits an MKPDU that contains a local MKA participant's member identifier (MI). A newly detected peer should start in the "P" state, and then transition to "L" in a matter of 2 to 4 seconds. A peer remaining in "P" indicates that the remote peer is not acknowledging the local peer's existence.
- **Connect Status**—Represents the controlled port state machine's "connect" variable. States are defined in *IEEE802.1X-2010 clause 12.3 CP state machine interfaces*:
 - **Pending**—Prevent connectivity by clearing the controlledPortEnabled parameter. Controlled port traffic is dropped.
 - **Authenticated**—Provide unsecured connectivity, setting controlledPortEnabled. Controlled port traffic is unencrypted.
 - **Secure**—Provide secure connectivity, using SAKs provided by the KaY (when available) and setting controlledPortEnabled when those keys are installed and in use, as specified in detail by the CP state machine. Controlled port traffic is encrypted.



Note

ExtremeXOS never chooses 'Unauthenticated' or 'Authenticated' access, but these options are allowed by the IEEE802.1X-2010 standard, so these cases may arise when interoperating with MKA/MACsec devices from other vendors.

- **Key Server**—Key server status:
 - **None**—Key server has yet to be elected (if persisting in this state, verify MACsec peer is enabled and PSKs are identical).
 - **Local**—This port has been elected key server.
 - **Peer**—Remote port has been elected key server.
- **MACsec**—Displays packet and byte statistics for both transmit and receive secure channels (SCs). Packet counters are 32-bits, while byte counters are 64-bits.
- **Usage**—Displays per-port MACsec usage information.

Example

The following example shows MKA and MACsec information for ports 25 and 50:



Note

To accommodate the width of the page, the MACsec columns are shown below the MKA content. In the actual output from the command, these columns appear beside each other.

```
# show macsec ports 25,50
MAC Security
-----MKA-----
      Local      Peer
Port  MACsec  Message  Message  Peer   Connect  Key
      Enabled  Number   Number  Status Status  Server
=====
25    Yes      0         -   N/A    PENDING  None
50    Yes     162244    162361  L      SECURE   Peer
=====

# show macsec ports 25,50
MAC Security
-----SecY-Tx-SC----SecY-Rx-SC-----
      Local      Peer
Encrypted  Octets  OK  Octets
Packets   Encrypted  Packets  Decrypted
=====+
      -         -         -         -
      1658      79584     2318     55827
=====
```

The following example shows MACsec usage on ports 1-2 and 49-56:

```
# show macsec ports 1-12,49-56 usage
Subject
Port      Link  to BW  MACsec  Allocated
      Speed  Maximum?  Enabled  Bandwidth
=====
1         1.0Gbps  Yes     Yes     1.0Gbps
2         1.0Gbps  Yes     Yes     1.0Gbps
3         1.0Gbps  Yes     Yes     1.0Gbps
4         1.0Gbps  Yes     Yes     1.0Gbps
5         1.0Gbps  Yes     Yes     1.0Gbps
6         1.0Gbps  Yes     Yes     1.0Gbps
7         1.0Gbps  Yes     Yes     1.0Gbps
8         1.0Gbps  Yes     Yes     1.0Gbps
9         1.0Gbps  Yes     Yes     1.0Gbps
10        1.0Gbps  Yes     Yes     1.0Gbps
11        1.0Gbps  Yes     No      -
12        1.0Gbps  Yes     No      -
49        10.0Gbps  Yes     No      -
50        10.0Gbps  Yes     No      -
51        10.0Gbps  Yes     No      -
52        10.0Gbps  Yes     No      -
53        10.0Gbps  Yes     Yes     10.0Gbps
55        10.0Gbps  Yes     No      -
```

History

This command was first available in ExtremeXOS 30.1.

The **usage** option was first available in ExtremeXOS 31.5

Platform Availability

This command is available on the following platforms.



Note

The MACsec feature requires the installation of the MAC Security feature pack license.

Platform	Ports
ExtremeSwitching 5320	All ports of all models except stacking ports.
ExtremeSwitching 5420	All ports of all models except stacking ports.
ExtremeSwitching 5520	All ports, except 5520-VIM-4X and 5520-24X 10G ports
ExtremeSwitching 5720	All ports of all models except stacking ports.

show macsec ports configuration

```
show macsec ports port-list configuration
```

Description

Displays a table of all configurable parameters.

Syntax Description

ports	Specifies ports to show configurable parameters on.
<i>port_list</i>	Lists which ports to show configurable parameters on.
configuration	Selects showing configurable MKA and MACsec parameters.

Default

N/A.

Usage Guidelines

None.

Example

The following example shows configurable parameters for ports 1 through 3:

```
# show macsec ports 1-3 configuration
MACsec      MKA Actor's  Cipher      Include  Replay Protect  MKA
Port        CA Name      Priority    Suite     SCI        Window Size  Lifetime
-----
1           BigCA        0x10      gcm-aes-128  Disabled          0
6
2           None         0x10      gcm-aes-128  Disabled          0
```

6						
3	BigCA	0x10	gcm-aes-128	Disabled	0	10

History

This command was first available in ExtremeXOS 30.1.

Cipher information added in ExtremeXOS 30.2.

MKA lifetime information was added in ExtremeXOS 31.5.

Platform Availability

This command is available on the following platforms.



Note

The MACsec feature requires the installation of the MAC Security feature pack license.

Platform	Ports
ExtremeSwitching 5320	All ports of all models except stacking ports.
ExtremeSwitching 5420	All ports of all models except stacking ports.
ExtremeSwitching 5520	All ports, except 5520-VIM-4X and 5520-24X 10G ports
ExtremeSwitching 5720	All ports of all models except stacking ports.

show macsec ports detail

```
show macsec ports port-list detail
```

Description

Displays configuration, status, and statistics for both MKA and MAC Security (MACsec).

Syntax Description

ports	Specifies ports to show MKA and MACsec detailed information on.
<i>port_list</i>	Lists which ports to show MKA and MACsec detailed information on.
detail	Selects showing detailed MACsec port information.

Default

N/A.

Example

The following example shows detailed MACsec information for port 25:

```
# show macsec ports 25 detail
PAE Port Table
-----
Port: 25
Port Capabilities          : 0x30
  Supplicant      : No
  Authenticator   : No
  MKA              : Yes
  MACsec          : Yes
  Announcements   : No
  Listener        : No
  Virtual Ports   : No
Virtual Ports Enable       : Disabled
Logon Enable               : Enabled
Authenticator Enable       : Enabled
Supplicant Enable         : Disabled
KaY MKA                   : Enabled
Announcer                : Disabled
Listener                 : Disabled

LOGON Table
-----
Connect          : SECURE
Port Valid       : True

NID Table
-----
UseEAP           : Never
UnauthAllowed    : Never
UnsecuredAllowed : mkaServer
UnauthenticatedAccess : noAccess
Access Capabilities : 0x08
  eap             : No
  eapMka          : No
  eapMkaMacSec    : No
  mka             : No
  mkaMacSec       : Yes
  vendorSpecific  : No

KaY MKA Table
-----
MKA Active          : True
MKA Authenticated   : False
MKA Secured         : True
MKA Failed          : False
MKA Actor SCI       : 00-04-96-99-39-93-00-19
MKA Actor's Priority : 0x2
MKA Life Time       : 10s
MKA Key Server SCI  : 00-04-96-99-39-93-00-19
MKA Key Server Priority : 0x2
MACsec Confidentiality Offset : 0
MACsec Desired      : True
MACsec Protect      : True
MACsec Replay Protect : True
MACsec Validate     : True
MACsec Protection
  Local MACsec Capability : Integrity, Confidentiality with Offset 0, 30, or 50
  Peer MACsec Capability  : Integrity, Confidentiality with Offset 0, 30, or 50
  Negotiated Protection   : Integrity, Confidentiality with Offset 0
MACsec Cipher Suite Admin : gcm-aes-256
MACsec Cipher Suite Oper  : gcm-aes-256
```

```

MKA Tx Key Number      : 6
MKA Tx Association Number : 1
MKA Rx Key Number      : 6
MKA Rx Association Number : 1

MKA Participant Table
-----
CA Name      : My256bitCA
CAK Name (CKN) : Switch1toSwitch2
Cached      : False
Active      : True
Retain      : False
ActivateControl : Default
Principal   : True

Potential Peer List :
Live Peer List      :
  MN, SCI : 26, 00-04-96-99-17-23-00-33
SecY Config Table
-----
Protect Frames:      Enabled
Validate Frames:     Strict
Replay Protect:      Enabled
Replay Protect Window: 0 frames
SecTAG Transmit Options
  Include SCI:        Disabled
  Use ES:              Disabled
  Use SCB:            Disabled

SecY Receive SA AN-1 Table
-----
State:              inUse
Next PN:            1899969
Created Time:       Fri Mar 22 10:55:16 2019

SecY Receive SC Table
-----
SCI:                00-04-96-99-17-23-00-33
State:              inUse
Current SA:         1
Created Time:       Fri Mar 22 10:55:16 2019

SecY Transmit SA AN-1 Table
-----
State:              inUse
Next PN:            1375880
Created Time:       Fri Mar 22 10:55:16 2019

SecY Transmit SC Table
-----
SCI:                00-04-96-99-39-93-00-19
State:              inUse
Encoding SA:        1
Enciphering SA:     0
Created Time:       Fri Mar 22 10:38:27 2019

SecY Interface Statistics
-----
SecY:
  Tx Untagged Pkts   : 0
  Tx Too Long Pkts   : 0
  Rx Untagged Pkts   : 0
  Rx No Tag Pkts     : 57046
  Rx Bad Tag Pkts    : 0

```

```

Rx Unknown SCI Pkts      : 0
Rx No SCI Pkts          : 0
Rx Overrun Pkts         : 0

Transmit:
  Secure Channel
    Protected Pkts       : 0
    Encrypted Pkts       : 4185922
    Octets Protected     : 0
    Octets Encrypted     : 6262129739

  Secure Association     : AN-1
    Protected Pkts       : 0
    Encrypted Pkts       : 4185922

Receive:
  Secure Channel, SCI: 00-04-96-99-17-23-00-33
    Late Pkts            : 0
    Not Valid Pkts       : 0
    Delayed Pkts         : 0
    Unchecked Pkts       : 0
    OK Pkts              : 1753184
    Octets Validated     : 0
    Octets Decrypted     : 2629771596

  Secure Association     : AN-1
    Not Valid SA Pkts    : 0
    OK Pkts              : 1753184

```

History

This command was first available in ExtremeXOS 30.1.

Cipher information was added in ExtremeXOS 30.2.

MKA lifetime information was added in ExtremeXOS 31.5.

Platform Availability

This command is available on the following platforms.



Note

The MACsec feature requires the installation of the MAC Security feature pack license.

Platform	Ports
ExtremeSwitching 5320	All ports of all models except stacking ports.
ExtremeSwitching 5420	All ports of all models except stacking ports.
ExtremeSwitching 5520	All ports, except 5520-VIM-4X and 5520-24X 10G ports
ExtremeSwitching 5720	All ports of all models except stacking ports.

show macsec ports usage

```
show macsec ports port-list usage
```

Description

Displays per-port MAC Security (MACsec) usage data in tabular format.

Syntax Description

ports	Specifies ports to show information on.
<i>port_list</i>	Lists which ports to view MACsec information on.
usage	Specifies to display per-port MACsec usage information.

Default

N/A.

Usage Guidelines

Displays per-port MACsec usage information. This command also displays the amount of bandwidth consumed by the enabled ports (assumes maximum throughput), as well as the maximum bandwidth allowed.

If the switch type has no port limit, then the "Port Count Limit" is displayed as N/A. If the switch has no bandwidth limit, then the "Bandwidth Maximum" is also displayed as N/A.



Note

The CLI prevents you from exceeding the port limit but not the bandwidth limit. If the instantaneous traffic across all MACsec-enabled ports exceeds the switch's bandwidth limit, then excess traffic will be dropped.

The following example shows MACsec usage on ports 1-2 and 49-56:

```
# show macsec ports 1-12,49-56 usage
      Subject
      Link   to BW   MACsec   Allocated
Port  Speed   Maximum? Enabled   Bandwidth
=====
1      1.0Gbps   Yes      Yes      1.0Gbps
2      1.0Gbps   Yes      Yes      1.0Gbps
3      1.0Gbps   Yes      Yes      1.0Gbps
4      1.0Gbps   Yes      Yes      1.0Gbps
5      1.0Gbps   Yes      Yes      1.0Gbps
6      1.0Gbps   Yes      Yes      1.0Gbps
7      1.0Gbps   Yes      Yes      1.0Gbps
8      1.0Gbps   Yes      Yes      1.0Gbps
9      1.0Gbps   Yes      Yes      1.0Gbps
10     1.0Gbps   Yes      Yes      1.0Gbps
11     1.0Gbps   Yes      No       -
12     1.0Gbps   Yes      No       -
49     10.0Gbps  Yes      No       -
50     10.0Gbps  Yes      No       -
51     10.0Gbps  Yes      No       -
52     10.0Gbps  Yes      No       -
53     10.0Gbps  Yes      Yes      10.0Gbps
55     10.0Gbps  Yes      No       -
```


History

This command was first available in ExtremeXOS 30.1.

The **usage** option was first available in ExtremeXOS 31.5

Platform Availability

This command is available on the following platforms.



Note

The MACsec feature requires the installation of the MAC Security feature pack license.

Platform	Ports
ExtremeSwitching 5320	All ports of all models except stacking ports.
ExtremeSwitching 5420	All ports of all models except stacking ports.
ExtremeSwitching 5520	All ports, except 5520-VIM-4X and 5520-24X 10G ports
ExtremeSwitching 5720	All ports of all models except stacking ports.

show macsec usage

```
show macsec usage
```

Description

Displays a per-port view of MAC Security (MACsec) usage information.

Syntax Description

usage	Specifies to display per-port MACsec usage information.
--------------	---

Default

N/A.

Usage Guidelines

Displays the number of ports that have MACsec enabled, and the maximum number of ports allowed, per slot. This command also displays the amount of bandwidth consumed by the enabled ports (assumes maximum throughput), as well as the maximum bandwidth allowed.

If the switch type has no port limit, then the "Port Count Limit" is displayed as N/A. If the switch has no bandwidth limit, then the "Bandwidth Maximum" is also displayed as N/A.



Note

The CLI prevents you from exceeding the port limit but not the bandwidth limit. If the instantaneous traffic across all MACsec-enabled ports exceeds the switch's bandwidth limit, then excess traffic will be dropped.

The following example shows MACsec usage information:

```
# show macsec usage
      ----Port Count----  -----Bandwidth-----
Slots      #Enabled      Limit  Allocated  Maximum
=====  =====
Slot-1          1         48    1.0Gbps   50.0Gbps
```

The following example shows MACsec usage information on ports 1-5:

```
# show macsec ports 1-5 usage
      Subject
Port      Link to BW  MACsec  Allocated
      Speed  Maximum?  Enabled  Bandwidth
=====  =====
1          1.0Gbps  Yes      No        -
2          1.0Gbps  Yes      No        -
3          1.0Gbps  Yes      No        -
4          1.0Gbps  Yes      No        -
5          1.0Gbps  Yes      No        -
```

History

This command was first available in ExtremeXOS 31.5.

Platform Availability

This command is available on the following platforms.



Note

The MACsec feature requires the installation of the MAC Security feature pack license.

Platform	Ports
ExtremeSwitching 5320	All ports of all models except stacking ports.
ExtremeSwitching 5420	All ports of all models except stacking ports.
ExtremeSwitching 5520	All ports, except 5520-VIM-4X and 5520-24X 10G ports
ExtremeSwitching 5720	All ports of all models except stacking ports.

show management

```
show management
```

Description

Displays the *SNMP* and CLI settings configured on the switch and the SNMP statistics.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

The following management output appears:

- Enable/disable state for Telnet, and SNMP access.
- Login statistics.
 - Enable/disable state for idle timeouts.
 - Maximum number of CLI sessions.
- SNMP community strings.
- SNMP trap receiver list.

For ExtremeXOS 11.0 and later, the following management output appears:

- SNMP trap receiver source IP address.
- SNMP statistics counter.
- SSH access state (enabled/disabled), idle time, rekey interval, and web access.
- CLI configuration logging.
- SNMP access states of v1, v2c disabled and v3 enabled.

If all three types of SNMP access are enabled or disabled, SNMP access is displayed as either Enabled or Disabled.

For ExtremeXOS 11.1 and later, the following management output appears:

- Enable/disable state for RMON.

For ExtremeXOS 11.2 and later, the following management output appears:

- Access-profile usage configured via *ACLs* for additional Telnet and SSH2 security.

For ExtremeXOS 11.6 and later, the following management output appears:

- CLI scripting settings
 - Enable/disable state.
 - Error message setting.
 - Persistence mode.

For ExtremeXOS 12.4 and later, the following management output appears:

- Dropped SNMP packet counter.

For ExtremeXOS 12.5 and later, the following management output appears:

- CLI prompting.
- SNMP INFORM.

For ExtremeXOS 22.1 and later, configured journal size (historical list of commands) appears.

For ExtremeXOS 22.3 and later, SSHv2 rekey interval time appears.

For ExtremeXOS 30.2 and later, CLI history expansion and port notation status appear.

For ExtremeXOS 30.3 and later, moved CLI commands information appears. The **CLI hidden moved-keywords** field shows:

- **Hidden**—moved commands do not appear in the CLI.
- **Displayed with Help**—moved commands appear with help text to direct you to the preferred new syntax.
- **Displayed with No Help**—moved commands appear without help text to direct you to the preferred new syntax.

The **CLI moved-keywords hidden release** field shows if the switch has been configured to hide moved commands, which version of ExtremeXOS version was running when the hide command was issued. To clear the ExtremeXOS version in this field, execute `unconfigure switch {all | erase [all | nvram]}` using the **all**, **erase all**, or **nvram** option. This also sets the moved command option back to **Displayed with Help**.

Example

The following command displays configured SNMP settings on a switch:

```
# show management
CLI idle timeout           : Enabled (20 minutes)
CLI max number of login attempts : 3
CLI max number of sessions  : 8
CLI paging                 : Enabled (this session only)
CLI space-completion       : Disabled (this session only)
CLI configuration logging   : Disabled (with expansion)
CLI journal size           : 100
CLI password prompting only : Disabled
CLI hidden moved-keywords   : Displayed with Help
CLI moved-keywords hidden release: 30.3
CLI RADIUS cmd authorize tokens : 2
CLI scripting              : Disabled (this session only)
CLI scripting error mode    : Ignore-Error (this session only)
CLI persistent mode        : Persistent (this session only)
CLI prompting              : Enabled (this session only)
CLI screen size            : 24 Lines 80 Columns (this session only)
CLI refresh                : Enabled
CLI history expansion       : Enabled
Current system port notation : port
Configured system port notation : port
Telnet access              : Enabled (tcp port 23 vr all)
                           : Access Profile : not set
SSH access                 : Enabled (Key valid, tcp port 22 vr all)
                           : Secure-Mode   : Off
                           : Access Profile : not set
SSH2 idle time             : 60 minutes
SSH2 rekey interval        : 4096 MB and no time limit
Web access                 : Enabled (tcp port 80)
```

```

: Access Profile : not set
Total Read Only Communities      : 1
Total Read Write Communities     : 1
RMON                             : Disabled
SNMP access                      : Enabled
                                : Access Profile : not set
SNMP Notifications               : Enabled
SNMP Notification Receivers      : None
SNMP stats:      InPkts 0      OutPkts 0      Errors 0      AuthErrors 0
                  Gets  0      GetNexts 0      Sets  0      Drops 0
SNMP traps:      Sent  0      AuthTraps Enabled
SNMP inform:     Sent  0      Retries 0      Failed 0

```

History

This command was first available in ExtremeXOS 10.1.

The trap receiver source IP address, SNMP counter statistics, SSH access, CLI logging, and SNMP access states were added to the output in ExtremeXOS 11.0.

The enabled/disabled state for RMON was added to the output in ExtremeXOS 11.1.

Additional Telnet and SSH2 information about ACL usage was added to the output in ExtremeXOS 11.2.

Information about CLI scripting including, the enabled/disabled state, error mode, and persistent mode was added to the output in ExtremeXOS 11.6.

The dropped SNMP packet counter (Drops) was added to the output in ExtremeXOS 12.4.

CLI prompting was added to the output in ExtremeXOS 12.5.

SNMP INFORM was added in ExtremeXOS 12.5.3.

Configured journal size (historical list of commands) was added in ExtremeXOS 22.1.

SSHv2 rekey interval time was added in ExtremeXOS 22.3.

Status of CLI configuration logging unabbreviated form (expansion) was added in ExtremeXOS 22.5.

CLI history expansion status and system port notation information were added in ExtremeXOS 30.2.

CLI moved commands information was added in ExtremeXOS 30.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show mcast cache

```

show mcast {ipv4 | ipv6} cache {{vlan} name} {{[group grpaddressMask |
  grpaddressMask] {source sourceIP | sourceIP}} {type [snooping | pim |
  mvr]} | {summary}}

```

Description

Displays multicast cache information.

The display can be limited to entries for specific VLANs or groups, and it can be limited to specific types of entries, such as those created by snooping protocols, PIM, or MVR.

Syntax Description

ipv4	Specifies the IPv4 address family.
ipv6	Specifies the IPv6 address family.
<i>name</i>	Specifies a <i>VLAN</i> name.
<i>grpaddressMask</i>	Specifies a multicast group address and mask.
<i>sourceIP</i>	Specifies the source IP address for a multicast group.
snooping	Limits the display to cache entries created by PIM or <i>IGMP</i> snooping.
pim	Limits the display to cache entries created by PIM.
mvr	Limits the display to cache entries created by MVR.
summary	Specifies the summary display format.

Default

Displays information for all entries in the multicast cache.

Usage Guidelines

If the `configure forwarding ipmc lookup-key mac-vlan` command is configured, the following message displays:

NOTE: Traffic is forwarded based on MAC address. Actual traffic forwarded based on the installed MAC address need not be the same displayed in this command, if overlapping IP multicast addresses are used in the network.

If the mode is **mixed-mode**, the following message displays:

NOTE: Traffic could be forwarded based on MAC address. Actual traffic forwarded based on the installed MAC address need not be the same displayed in this command, if overlapping IP multicast addresses are used in the network.

Example

The following command displays all multicast cache information:

```
# show mcast cache
Snooping/MVR Cache Timeout: 300 sec
Type Group          Sender          Age  InVlan
snoop 225.1.1.1      222.222.222.222  17   snvlan
Vlan          Port      Vid
snvlan        2         400
23            400
snoop 224.0.0.5      100.1.2.2       2    pmvlan2
Vlan          Port      Vid
pmvlan2       4         402
```

```
snoop 224.0.0.5          100.1.3.3          17  pmvlan3
Vlan                    Port      Vid
pmvlan3                 23      403
snoop 224.0.0.13        100.1.2.2          11  pmvlan2
Vlan                    Port      Vid
pmvlan2                 4       402
snoop 224.0.0.13        100.1.3.3          14  pmvlan3
Vlan                    Port      Vid
pmvlan3                 23      403
pim 226.1.1.1           100.1.1.12         0   pmvlan1
Vlan                    Port      Vid
pmvlan2                 4       402
pmvlan3                 23      403
Multicast cache distribution:
5 entries from Snooping          0 entries from MVR          1 entries from PIM
Total Cache Entries: 6
```

The following command displays summary cache information for VLAN pmvlan1:

```
# show mcast cache vlan pmvlan1 summary
Snooping/MVR Cache Timeout: 300 sec
=====MULTICAST CACHE SUMMARY=====
Multicast cache distribution:
5 entries from Snooping          0 entries from MVR          1 entries from PIM
pmvlan1: Multicast cache distribution:
0 entries from Snooping          0 entries from MVR          1 entries from PIM
Total Cache Entries: 6
Total Cache Entries for VLAN pmvlan1: 1
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show mcast ipv6 cache

```
show mcast ipv6 cache {{vlan} name} {[group v6GrpAddressMask |
v6GrpAddressMask] {source v6SourceIP | v6SourceIP}} {type [snooping |
pim]} {with-in-port} | {summary}}
```

Description

Displays multicast cache information. The display can be limited to entries for specific VLANs or groups, and it can be limited to specific types of entries, such as those created by snooping protocols, or PIM.

Syntax Description

<i>name</i>	Specifies a VLAN name.
<i>v6GrpAddressMask</i>	Specifies a multicast group address and mask.
<i>v6SourceIP</i>	Specifies the source IP address for a multicast group.

snooping	Limits the display to cache entries created by MLD snooping.
pim	Limits the display to cache entries created by PIM.
summary	Specifies the summary display format.

Default

Displays information for all entries in the multicast cache.

Usage Guidelines

None.

Example

The following command displays all multicast cache information:

```
# show mcast ipv6 cache
Snooping Cache Timeout: 300 sec
(ff03::1 3001::1)
Type: snoop Age: 9 Ingress Vlan: v1
Vlan      Port      Vid
v1         25         4084
(ff03::1 3001::2)
Type: snoop Age: 9 Ingress Vlan: v1
Vlan      Port      Vid
v1         25         4084
Multicast cache distribution:
2 entries from Snooping          0 entries from PIM
Total Cache Entries: 2
```

The following command displays summary cache information for VLAN v1:

```
# show mcast ipv6 cache vlan v1 summary
Snooping Cache Timeout: 300 sec
=====MULTICAST CACHE SUMMARY=====
Multicast cache distribution:
2 entries from Snooping          0 entries from PIM
v1: Multicast cache distribution:
2 entries from Snooping          0 entries from PIM
Total Cache Entries: 2
Total Cache Entries for VLAN v1: 2
#
```

History

This command was first available in ExtremeXOS 15.2.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MLD snooping feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

show memory

```
show memory {slot slotid}
```

Description

Displays the current system memory information.

Syntax Description

slotid	Specifies slot number for the node in a stack. The value can be from 1 to 8.
---------------	--

Default

N/A.

Usage Guidelines

Viewing statistics on a regular basis allows you to see how well your network is performing. When you keep simple daily records, you 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.

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

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

You can also use the `show memory process name {slotslotid}` command to view the system memory and the memory used by the individual processes.

SummitStack Only

When you issue the command with out any parameters:

- From the stack manager or backup node, the stack displays the current system memory information for the master node and the back-up node in the Active Topology.
- From a standby node, the stack displays the current system memory information for the master node and the standby node in the Active Topology.

Reading the Output

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

- System memory information (both total and free).
- Current memory used by the individual processes.

The current memory statistics for the individual process also includes the name of the process.

In general, the free memory count for a switch decreases when one or more running processes experiences an increase in memory usage.

If you observe a continuous decrease in the free memory over an extended period of time, and you have not altered your switch configuration, please contact Extreme Networks Technical Support.

Example

The following command displays current system memory information for a switch:

```
show memory
```

The following is sample output from this command:

```
System Memory Information
-----
Total DRAM (KB): 262144
System      (KB): 25852
User        (KB): 96608
Free        (KB): 139684
Memory Utilization Statistics
-----
Process Name      Memory (KB)
-----
aaa                13468
acl                11420
bgp                 0
cfgmgr             8336
cli               41040
cna                 0
devmgr             8452
dirser             7068
dosprotect         8264
eaps              18784
edp                9780
elrp              10040
ems               10672
epm               15520
esrp              16728
etmon             18924
exacl              30
exdos              8
exfib              3
exosmc             29
exosnvram          4
exosq              36
exsflow           10
exsnoop           20
exsshd            9272
exvlan            290
fdb              12908
hal               64768
lldp              8816
mcmgr             17836
msgsrv            6960
netLogin          8924
netTools          11524
nettx              70
nodemgr           9636
ospf              18124
ospfv3             0
pim               15996
```

```

poe                8936
polMgr             7576
rip                17736
ripng              0
rtmgr             16016
snmpMaster         15416
snmpSubagent       26428
stp                10768
telnetd            8464
tftpd              7584
thttpd            11344
vlan               9660
vrrp               11184
xmld               9148

```

The following command displays current system memory information for a stack, where slot 1 is the master and slot 6 is the backup:

```

Slot-1 stacK.3 # show memory
System Memory Information
-----
Slot-1    Total DRAM (KB): 262144
Slot-1    System      (KB): 25476
Slot-1    User        (KB): 132256
Slot-1    Free        (KB): 104412
Slot-6    Total DRAM (KB): 262144
Slot-6    System      (KB): 25476
Slot-6    User        (KB): 122820
Slot-6    Free        (KB): 113848
Memory Utilization Statistics
-----
Card Slot Process Name      Memory (KB)
-----
Slot-1 1   aaa              2548
Slot-1 1   acl              2960
Slot-1 1   bgp              0
Slot-1 1   brm              2428
Slot-1 1   cfgmgr           3256
Slot-1 1   cli              16932
Slot-1 1   devmgr           2708
Slot-1 1   dirser           1916
Slot-1 1   dosprotect       1972
Slot-1 1   eaps             6976
Slot-1 1   edp              2656
Slot-1 1   elrp             2640
Slot-1 1   elsm             2592
Slot-1 1   ems              2764
Slot-1 1   epm              3092
Slot-1 1   esrp             2844
Slot-1 1   etmon            16264
...
Slot-6 6   aaa              2440
Slot-6 6   acl              2872
Slot-6 6   bgp              0
Slot-6 6   brm              2396
Slot-6 6   cfgmgr           2776
Slot-6 6   cli              16292
Slot-6 6   devmgr           2672
Slot-6 6   dirser           1836
Slot-6 6   dosprotect       1944
Slot-6 6   eaps             6924
Slot-6 6   edp              2624
Slot-6 6   elrp             2628
Slot-6 6   elsm             2564

```

```
Slot-6 6   ems      2744
Slot-6 6   epm      2976
Slot-6 6   esrp     2792
Slot-6 6   etmon    10068
...
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show memory process

```
show memory process name {slot slotid}
```

Description

Displays the current system memory and that of the specified process.

Syntax Description

<i>name</i>	Specifies the name of the process.
<i>slotid</i>	In a SummitStack, <i>slotid</i> specifies the slot number of the node in the stack topology. The value can be from 1 to 8.

Default

N/A.

Usage Guidelines

Viewing statistics on a regular basis allows you to see how well your network is performing. When you keep simple daily records, you 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.

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

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

You can also use the `show memory {slot slotid}` command to view the system memory and the memory used by the individual processes.

SummitStack Only

When you issue the command with out any parameters:

- From the stack manager or backup node, the stack displays current system memory and that of the specified process running on the master node and the back-up node in the Active Topology.
- From a standby node, the stack displays current system memory and that of the specified process running on the master node and the standby node in the Active Topology.

Reading the Output

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

- System memory information (both total and free).
- Current memory used by the individual processes.

The current memory statistics for the individual process also includes the following:

- For SummitStacks, the slot number.
- The name of the process.

Example

The following example displays system memory and VRRP memory usage:

```
show memory process vrrp
```

The following is sample output from a Summit switch:

```
System Memory Information
-----
Total DRAM (KB): 262144
System      (KB): 25852
User        (KB): 96608
Free        (KB): 139684
Memory Utilization Statistics
-----
Process Name      Memory (KB)
-----
vrrp               11184
```

The following is sample output from a SummitStack:

```
Slot-1 stacK.4 # show memory process "aaa"
System Memory Information
-----
Slot-1    Total DRAM (KB): 262144
Slot-1    System      (KB): 25476
Slot-1    User        (KB): 132276
Slot-1    Free        (KB): 104392
Slot-6    Total DRAM (KB): 262144
Slot-6    System      (KB): 25476
Slot-6    User        (KB): 122820
Slot-6    Free        (KB): 113848
```

Memory Utilization Statistics

```

-----
Card Slot Process Name      Memory (KB)
-----
Slot-1 1   aaa              2548
Slot-6 6   aaa              2440

```

History

This command was first available in an ExtremeXOS 11.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show meter

```
show meter meter_name {ports [port_group | port_list]} {out-actions}
```

Description

Displays the configured meters.

Syntax Description

<i>meter_name</i>	Specifies the meter name.
ports	Ports to display.
<i>port_group</i>	Port group name.
<i>port_list</i>	Port list separated by a comma.
out-actions	Out-of-profile actions.

Default

N/A.

Usage Guidelines

None.

Example

The following sample output displays meters on the switch:

```

#show meter

-----
Name      Committed Rate (Kbps)  Peak Rate (Kbps)

```

```
-----
peggy          1000000--
```

The following example uses the *port-group* variable:

```
# show meter ports ingMeterGrpA

Group: ingMeterGrpA
Ports: 1-3
Name          Committed Rate      Peak Rate  Burst Size (Kb)
-----
ingmeter2     21000 Kbps      30000 Kbps      2
ingmeter3     30000 Pps       40000 Pps       5
ingmeter4     21000 Kbps      --              2
ingmeter5     30000 Pps       --              5
```

The following example uses *port_group* **out-actions**:

```
# show meter ingmeter2 ports ingMeterGrpB out-actions

Group: ingMeterGrpB
Ports: 4-5

Out-actions
Meter          Log  Trap  Disable  Drop  Drop  Drop
Name           --  ---  Port    Port  Precedence
-----
ingmeter2      --  On   On       On   63
```

The following example uses *port_list*:

```
# show meter ports 1-4

Name          Committed Rate      Peak Rate  Burst Size(Kb)  Port
-----
ingmeter2     21000 Kbps      30000 Kbps      2 1
ingmeter5     30000 Kbps      --           5 1
ingmeter2     21000 Kbps      --           2 2
ingmeter5     30000 Kbps      --           5 2
ingmeter2     21000 Kbps      --           2 3
ingmeter5     30000 Kbps      --           5 3
ingmeter5     2000 Pps       10000 Pps      3 4
ingmeter7     10000 Kbps      --           5 4
```

The following example uses *port_list* **out_actions**:

```
# show meter "ingmeter2" ports 4-6 out-actions

Out-
actions
Meter          Log  Trap  Disable  Drop  Drop  Drop  Port
Name           --  ---  Port    Port  Precedence (dscp)
-----
ingmeter2      On  On   On       --   63  4
```

ingmeter2	On	On	On	--	63	5
ingmeter2	On	On	On	On	--	6



Note

You configure a peak rate for QoS meters using the `configure meter metername {committed-rate cir [Gbps | Mbps | Kbps | Pps]} {max-burst-size burst-size [Kb | Mb | packets]} {out-actions [{disable-port} {drop | set-drop-precedence {dscp [none | dscp-value]} {dot1p [none | dot1p-value]}]}` command.

History

This command was first available in ExtremeXOS 11.4.

The **ports**, *port_group*, *port_list*, and **out-actions** options were added in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show meter out-of-profile

```
show meter {metername} out-of-profile [{disabled-ports} ports [portlist
| port_group] | global-count}
```

Description

This command displays the meters that are out-of-profile and the ports that have been disabled as a result of a meter that is out-of-profile. The show screen will display the out-of-profile status and rate-limit counters for both per-port and global meters. Both the per-port and global meter counters and status are an aggregate of the rule based counters and status which includes both *ACL* and dot1p rules. In addition a screen will display the global out-of-profile count for per-port meters. The global count is used for debugging systems that do not support per-port counters in the hardware. This global counter for per-port meters is an aggregate of the rule based counters which includes both ACL and dot1p rules.

Syntax Description

<i>metername</i>	Meter name.
disabled-ports	Show the meter out-of-profile status that resulted in disable-port action.
ports	Show the meter applied to a specified port list.
<i>portlist</i>	Port list separated by a comma or -
<i>port_group</i>	Port group name.
global-count	Counter of all the rules (ACL and dot1p) using a per-port meter.

Default

N/A.

Usage Guidelines

None.

Example

```
#show meter out-of-profile ports ingMeterGrp

Group: ingMeterGrp
Ports: 4-5

Name              Status      Disabled  Rate-Limit Counter  Port
-----
ingmeter1         Out of profile  Yes      1234567             4
ingmeter1         Ok              No        0                    5
ingmeter2         Out of profile  No        0                    4
ingmeter2         Out of profile  No        1467                 5
ingmeter3         Out of profile  No        0                    4
ingmeter3         Ok              No        0                    5
```

History

This command was first available in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show mirror

```
show mirror [mirror_name | control_index | mirror_name_li] | [all |
enabled]
```

Description

Displays various show output for mirror instances.

Syntax Description

<i>mirror_name</i>	Displays the mirror name.
<i>mirror_name_li</i>	Displays mirror instance name for Lawful Intercept account only. You must have lawful intercept user privileges to specify this variable.
all	Displays all mirror instances.
enabled	Displays only enabled mirror instances.
<i>control_index</i>	Selects the particular mirror MIB instance, as defined by a control index (1-4), to show information about.

Default

N/A.

Usage Guidelines

Use this command to display mirror statistics and determine if mirroring is enabled or disabled on the switch.

For mirror instances to remote IP addresses, one the following statuses appears:

- **Disabled**—User has disabled the mirror instance.
- **Down. No route**—No route is available to reach the remote IP address.
- **Down. IP forwarding not enabled on VLAN to route**—Route is available to reach the remote IP address, but egress VLAN does not have IP forwarding enabled (see [enable ipforwarding](#) on page 2148).
- **Down. Ping timed out**—Route is available, IP address forwarding is enabled on the egress VLAN, but ping check is “on” and a timely ping response from the remote IP address was not received during three consecutive ping requests, sent every five seconds.
- **Up. Active**—Remote IP address is reachable and ARP entry is resolved. Ping health check is either “off”, or is “on” and a timely ping response has been received. The status is “Up. Active” for the remote IP address in “Up” state with the highest configured priority, and therefore is sent mirrored packets.
- **Up. Standby**—Remote IP address is reachable and ARP entry is resolved. Ping health check is either “off”, or is “on” and a timely ping response has been received. Status is “Up. Standby” for remote IP addresses in “Up” state, but whose configured priority is below the priority of the remote IP address that is “Up. Active”.



Note

The mirror instance for a remote IP address is only installed in hardware if the status is “Up.Active”.

Example

The following command displays switch mirroring information:

```
# show mirror

DefaultMirror    (Disabled)    Description:    Default Mirror Instance, created
automatically
  Mirror to port: -

PolicyMirror-1003 (Enabled)
  Mirror Destination Indices: 1 - Enabled
  Description:
  Mirror to port: 3

TunMir    (Enabled)
  Mirror Destination Indices: 3 - Enabled
  Description:
  Mirror to remote IP: 192.168.97.100    VR          : VR-Default
  From IP          : 192.168.97.184    Ping check: On
  Status           : Up
```

```

Mirrors defined:          3
Mirrors enabled:          2 (Maximum 4)
HW filter instances used: 0 (Maximum 128)
HW mirror instances used: 0 ingress, 0 egress (Maximum 4 total, 2 egress)
Remote IP protocol type:  Trans Ether Bridging (0x6558)

Destination Control Index 1: Enabled (Active)
Destination Control Index 3: Enabled (Active)

```

The following example displays output for a lawful intercept user session:

```

* show mirror
DefaultMirror    (Disabled)
  Description:    Default Mirror Instance, created automatically
  Mirror to port: -

law_mirror       (Enabled)
  Description:    user for lawful intercept
  Mirror to port: 3
  Source filter instances used : 2
    Port 7, all vlans, ingress only
    Port 8, all vlans, ingress only

main_mirror      (Enabled)
  Description:
    Mirror instance for Admin
  Mirror to port: 2
  Source filter instances used : 1
    Port 10, all vlans, ingress only

Mirrors defined:          3
Mirrors enabled:          2 (0 with egress filters)
HW filter instances used: 3 (Maximum 128)
Remote IP protocol type:  ERSPAN v1 (0x88BE)

```

The following example shows mirroring to remote IP addresses:

```

# show mirror

analytics_chicago_1  (Enabled)
  Description:
    Mirror to remote IP: 5.1.1.2          VR          : VR-Default
    From IP              : Auto source IP   Ping check: On
    Priority              : 50
    Status                : Up. Active

    Mirror to remote IP: 4.1.1.2          VR          : VR-Default
    From IP              : Auto source IP   Ping check: On
    Priority              : 40
    Status                : Up. Standby

    Mirror to remote IP: 3.1.1.2          VR          : VR-Default
    From IP              : Auto source IP   Ping check: On
    Priority              : 30
    Status                : Down. Ping timed out

    Mirror to remote IP: 2.1.1.2          VR          : VR-Default
    From IP              : Auto source IP   Ping check: On
    Priority              : 20
    Status                : Up. Standby
    Source filter instances used : 1
      All ports, vlan v1, ingress only

analytics_seattle_2   (Enabled)

```

```

Description:
Mirror to remote IP: 5.6.7.8          VR          : VR-Default
From IP          : 10.1.1.1          Ping check: Off
Priority         : 50
Status          : Down. No route
Source filter instances used : 1
                Port 2, vlan v2, ingress only

bldg_1_sniffer   (Enabled)
Description:
Mirror to port: 5
Source filter instances used : 1
                All ports, vlan v3, ingress only

DefaultMirror    (Disabled)
Description:      Default Mirror Instance, created automatically
Mirror to port: -

Mirrors defined:      4
Mirrors enabled:      3 (Maximum 4)
HW filter instances used: 3 (Maximum 128)
HW mirror instances used: 3 ingress, 0 egress (Maximum 4 total, 2 egress)
Remote IP protocol type: ERSPAN v1 (0x88BE)

```

The following example shows information about Mirror MIB instance for control index "1":

```

# show mirror 1
Destination Control Index 1: Enabled (Active)

```

History

This command was first available in ExtremeXOS 15.3.

The *mirror_name_li* variable was added in ExtremeXOS 15.3.2.

Mirroring to remote IP addresses information was added in ExtremeXOS 22.4.

Mirror MIB instance (control index) information for policy-based mirrors was added in ExtremeXOS 30.2.

Redundant remote IP address information was added in ExtremeXOS 30.4.

GRE protocol type for mirror-to-remote IP addresses was added in ExtremeXOS 30.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show mlag peer

```
show mlag peer {peer_name}
```

Description

Displays information about an .

Syntax Description

<i>peer_name</i>	Specifies an alpha numeric string identifying the MLAG peer switch.
------------------	---

Default

N/A.

Usage Guidelines

Use this command to display configured items, MLAG peer switch state, MLAG group count, and health-check statistics.

Example

The following command displays information for an MLAG peer switch:

```
# show mlag peer
```

Following is sample output for the command:

Multi-switch Link Aggregation Peers:

```
MLAG Peer      : leftBD8k
VLAN           : isc          Virtual Router    : VR-Default
Local IP Address : 1.1.1.2     Peer IP Address  : 1.1.1.1
MLAG ports      : 2           Tx-Interval     : 1000 ms
Checkpoint Status : Up        Peer Tx-Interval : 1000 ms
Rx-Hellos       : 184         Tx-Hellos       : 184
Rx-Checkpoint Msgs : 12       Tx-Checkpoint Msgs : 12
Rx-Hello Errors  : 0          Tx-Hello Errors  : 0
Hello Timeouts   : 1          Checkpoint Errors : 0
Up Time          : 0d:0h:0m:10s Peer Conn.Failures : 1
Local MAC :00:04:96:11:22:44 Peer MAC :00:04:96:11:22:33
Config'd LACP MAC :None       Current LACP MAC :00:04:96:11:22:33
Multi-switch Link Aggregation Peers:

MLAG Peer      : rightBD8k
VLAN           : isc          Virtual Router    : VR-Default
Local IP Address : 1.1.1.1     Peer IP Address  : 1.1.1.2
MLAG ports      : 2           Tx-Interval     : 1000 ms
Checkpoint Statu : Up        Peer Tx-Interval : 1000 ms
Rx-Hellos       : 167         Tx-Hellos       : 167
Rx-Checkpoint Msgs: 12       Tx-Checkpoint Msgs : 12
Rx-Hello Errors  : 0          Tx-Hello Errors  : 0
Hello Timeouts   : 1          Checkpoint Errors : 0
Up Time          : 0d:0h:0m:7s Peer Conn.Failures : 1
Local MAC :00:04:96:11:22:44 Peer MAC :00:04:96:11:22:33
Config'd LACP MAC :None       Current LACP MAC :00:04:96:11:22:33
```

Following is sample output when an MLAG peer has been created but the IP address is yet to be configured:

```
* switch # show mlag peer switch101
```

Multi-switch Link Aggregation Peers:

```
MLAG peer      : switch101
VLAN           :          Virtual Router    :
Local IP address :          Peer IP address  :
```

```

MLAG groups      : 0          Tx-Interval      : N/A
Checkpoint Status : Down      Peer Tx-Interval : N/A
Rx-Hellos        : 0          Tx-Hellos        : 0
Rx-Checkpoint Msgs : 0          Tx-Checkpoint Msgs : 0
Rx-Hello Errors   : 0          Tx-Hello Errors   : 0
Hello Timeouts    : 0          Checkpoint Errors  : 0
Up Time: N/A      : 0          Peer Conn.Failures : 0
Local MAC         :00:04:96:11:22:44 Peer MAC      :00:04:96:11:22:33
Config'd LACP MAC :None        Current LACP MAC: 00:04:96:11:22:33

```

Following is sample output displaying LACP MAC for an MLAG peer:

```

* switch # show mlag peer switch101

Multi-switch Link Aggregation Peers:

MLAG Peer      :S2
VLAN           :isc          Virtual Router    :VR-Default
Local IP Address :1.1.1.1     Peer IP Address   :1.1.1.2
MLAG ports      :0           Tx-Interval       :1000 ms
Checkpoint Status :Up         Peer Tx-Interval   :1000 ms
Rx-Hellos       :153379      Tx-Hellos         :153895
Rx-Checkpoint Msgs :6         Tx-Checkpoint Msgs :14
Rx-Hello Errors  :0          Tx-Hello Errors   :0
Hello Timeouts   :0          Checkpoint Errors  :0
Up Time          :1d:17h:45m:8s Peer Conn. Failures :0
Local MAC        :00:04:96:11:22:44 Peer MAC          :00:04:96:11:22:33
Config'd LACP MAC :None        Current LACP MAC   :
:00:04:96:11:22:33            Config'd LACP MAC

```

Configured LACP MAC is the configured LACP MAC using the:

```
configure {mlag peer} peer_name lacp-mac lacp_mac_address
```

command.

If no MAC is configured,

```
Config'd LACP MAC
```

is shown as None.

If same LACP MAC is configured on both the switches, the current LACP MAC will be the same as

```
Config'd LACP MAC.
```

If LACP MAC is not configured on any of the MLAG peers or if a different MAC is configured on the peers,

```
Current LACP MAC
```

is different from Config'd LACP MAC and is selected from Local MAC/Peer MAC combination.

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, 5720 series switches.

show mlag ports

```
show mlag ports {port_list}
```

Description

Displays information about each group.

Syntax Description

<code>port_list</code>	Specifies one or more ports.
------------------------	------------------------------

Default

N/A.

Usage Guidelines

Use this command to display information about each MLAG group including local port number, local port status, remote MLAG port state, MLAG peer name, MLAG peer status, local port failure count, remote MLAG port failure count.

Local and remote link state and fail counts reflect the status of the entire LAG when a LAG is used in conjunction with an MLAG. For example, if 1 and 2 ports in a local LAG on the switch associated with an MLAG is down, the local link state will still show as ready and the associated local fail count will be incremented. The remote fail count shown at MLAG neighboring switch will also be incremented.

Example

The following command displays information for an MLAG group:

```
# show mlag ports
Local
MLAG      Local  Link   Remote      Local  Remote
Id        Port   State  Link   Peer      Status  Fail    Fail
=====
2         1:1    A      Up      leftBD8K   Up       0       0
1         1:2    A      Up      leftBD8K   Up       0       0
=====
Local Link State: A - Active, D - Disabled, R - Ready, NP - Port not present
Remote Link      : Up - One or more links are active on the remote switch,
Down - No links are active on the remote switch,
N/A - The peer has not communicated link state for this MLAG
port,
Virtual - MLAG peer switch does not have physical port.
Number of Multi-switch Link Aggregation Groups : 2
Convergence control                          : Fast
Reload Delay Interval                        : 30 seconds
Reload Delay                                : Disabled
Link Up Isolation                            : Off
```

The following command displays information about an MLAG group on ports 1 and 2:

```
show mlag port 1,2
Local
MLAG      Local  Link   Remote      Local  Remote
Id        Port   State  Link   Peer      Status  Fail    Fail
=====
100       1       A      Up      switch101  Up       0       2
101       2       A      Down    switch101  Up       0       1
```

```
=====
Local Link State: A - Active, D - Disabled, R - Ready, NP - Port not present
Remote Link: Up - One or more links are active on the remote switch,
Down - No links are active on the remote switch,
N/A - The peer has not communicated link state for this MLAG
group
Virtual - MLAG peer switch does not have physical port.

Number of Multi-switch Link Aggregation Groups: 2
Convergence Control                : Conserve Access Lists
Reload Delay Interval              : 30 seconds
Reload Delay                      : Disabled
Link Up Isolation                  : Off
```

History

This command was first available in ExtremeXOS 12.5.

MLAG reload delay timer information was added in ExtremeXOS 22.3.

Virtual port information was added in ExtremeXOS 22.7.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, 5720 series switches.

show mld

```
show mld {vlan} {name}
```

Description

This command can be used to display an MLD-related configuration and group information, per VLAN or for the switch as a whole.

Syntax Description

<i>name</i>	Specifies a VLAN name.
-------------	------------------------

Default

N/A.

Usage Guidelines

If you do not specify a VLAN, the command displays the switch configuration.

Example

The following is sample output from the `show mld` command:

```
# show mld
VLAN          IP Address          Flags  nLRMA  nLeMA  MLDver
Default       ::/0                ---iz-    0      0      0
v1            ::/0                U--iz-    0      5      0
Flags:  (f) Forwarding Enabled, (g) Fast-learning on, (i) MLD Enabled,
(m) Multicast Forwarding Enabled, (U) Interface Up,
(z) MLD Snooping Enabled.
(nLeMA) Number of Learned Multicast Addresses
(nLRMA) Number of Locally Registered Multicast Addresses
```

The following command displays the MLD configuration for VLAN v1:

```
# show mld v1
Query Interval      : 125 sec
Max Response Time  : 10 sec
Last Member Query   : 1 sec
Robustness          : 2
Interface on VLAN v1 is enabled and up.
inet6 ::/0
Locally registered multicast addresses:
Learned multicast addresses (Last Querier=fe80::204:96ff:fe3a:ce50):
ff02::2                ff02::1:ff56:5c2b
ff02::1:ff00:2          ff02::1:ff3a:ce50
ff02::1:ff55:5c27
s = static MLD member
Flags:
IP Fwding  NO      IPmc Fwding  NO      MLD YES
MLD Ver   v0      Snooping YES
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 multicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

show mld counters

```
show mld counters {{vlan} name}
```

Description

Use this command to display an MLD packet statistics.

Syntax Description

<i>name</i>	Specifies a <u>VLAN</u> name.
-------------	-------------------------------

Default

N/A.

Usage Guidelines

The following command displays the MLD configuration:

```
* topleft.74 # show mld counters
MLD Message type           Received   Originated   Forwarded
MLD Query (v1/v2)           0         20           0
MLDv1 Report                499        0           157
MLDv1 Done                  101        0           91
MLDv2 Report                0          0           0
Global Statistics:
MLD Packet unknown          0
MLD Packet Error            617
```

History

This command was first available in ExtremeXOS 15.2.

Platform Availability

This command is available on the platforms listed for the IPv6 multicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

show mld group

```
show mld group {{vlan} {name} | {v6grpipaddress}} {MLDv2}
```

Description

Lists the MLD group membership for the specified VLAN or group.

Syntax Description

<i>name</i>	Specifies a VLAN name.
<i>v6grpipaddress</i>	Specifies a group IPv6 address.
MLDv2	DisplayS the MLD group in MLDv2 format (if group record is MLDv2 compatible, otherwise displayS in earlier format). This option is not supported in this release.

Default

MLDv1.

Usage Guidelines

If no VLAN is specified, all VLANs are displayed. You can also filter the display by group address or by multicast stream sender address.

Example

The following command lists the MLD group membership for the VLAN accounting:

```
show mld group vtest3
```

Output from this command looks similar to the following:

Group Address	Ver	Vlan	Port	Age
ff03::1:1	1	vtest3	4:5	25
ff03::1:2	1	vtest3	4:5	25
ff02::1:ff22:124	1	vtest3	4:45	26
ff05::a:abcd	1	vtest3	4:15	23
ff05::a:abce	1	vtest3	4:15	23
ff02::1:ff22:112	1	vtest3	4:45	26
ff02::1:ff1f:a418	1	vtest3	4:45	26

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 multicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

show mld snooping

```
show mld snooping {vlan name | detail} {MLDv2}
```



Note
MLD snooping is not supported in this software release.

Description

Displays MLD snooping registration information and a summary of all MLD timers and states.

Syntax Description

<i>name</i>	Specifies a <u>VLAN</u> name.
detail	Displays the information in detailed format.
MLDv2	Displays the MLD group in MLDv2 format (if group record is MLDv2 compatible, otherwise displays in earlier format). This option is not supported in this release.

Default

MLDv1.

Usage Guidelines

The two types of MLD snooping entries are sender entry and subscribed entry.

The following information is displayed in a per-interface format:

- Group membership information
- Router entry
- Timeout information
- Sender entry

Example

Here is an example of the show output:

```
# show mld snooping
MLD Snooping Flood-list : none
MLD Snooping Proxy      : Enable
MLD Snooping Filters    : per-port
Vlan                    Vid   Port   #Senders #Receivers Router Enable
-----
Default                1      0      0         0         Yes
v1                      4084   0      0         0         Yes
25                      1      Yes
41                      2      No
42                      2      No
```

The following command displays MLD snooping registration information for the VLAN V1:

```
# show mld snooping v1
Router Timeout          :    260 sec
Host Timeout           :    260 sec
MLD Snooping Fast Leave Time : 1000 ms
VLAN v1                (4084) Snooping=Enabled
Port   Host                                     Age
Subscribed                                     Join Limit
25     fe80::204:96ff:fe3a:ce50                  13
ff02::1:ff3a:ce50                             No Limit
25     fe80::204:96ff:fe3a:ce50                  14
All Groups                                     No Limit
41     fe80::200:8ff:fe55:5c27                   13
ff02::1:ff00:2                                No Limit
41     fe80::200:8ff:fe55:5c27                   13
ff02::1:ff55:5c27                             No Limit
42     fe80::200:8ff:fe56:5c2b                   14
ff02::1:ff00:2                                No Limit
42     fe80::200:8ff:fe56:5c2b                   13
ff02::1:ff56:5c2b                             No Limit
s = static MLD member
#
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 multicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

show mld snooping vlan filter

```
show mld snooping {vlan} name filter
```

Description

Displays MLD snooping filters..

Syntax Description

<i>name</i>	Specifies a <i>VLAN</i> name.
-------------	-------------------------------

Default

None.

Usage Guidelines

Use this command to display MLD snooping filters configured on the specified VLAN. When no VLAN is specified, all the filters are displayed.

Example

The following command displays the MLD snooping filter configured on VLAN vlan101:

```
# show mld snooping vlan101 filter
Filter Port Flags
mldpermit0 5:10 a
Flags: (a) Active
```

History

This command was first available in ExtremeXOS 15.2.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MLD snooping feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

show mld snooping vlan static

```
show mld snooping vlan name static [group | router]
```

Description

Displays static MLD snooping entries.

Syntax Description

<i>name</i>	Specifies a <u>VLAN</u> name.
-------------	-------------------------------

Default

None.

Usage Guidelines

Use this command to display the MLD snooping static groups or router ports configured on the specified VLAN. When no VLAN is specified, all the static groups or router ports are displayed.

Example

The following command displays the MLD snooping static groups configured on VLAN vlan101:

```
show mld snooping vlan101 static group
```

The following is sample output for this command:

Group	Port	Flags
ff03::1:1:1	7	sa
ff03::1:1:1	15	sa
Flags: (s) Static, (a) Active		

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 multicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

show mld ssm-map

```
show mld ssm-map {v6groupnetmask} [{vr} vr_name]
```

Description

Displays the status of MLD-SSM mapping feature on a VR (if it is enabled or disabled), and displays the MLD-SSM mapping entries.

Syntax Description

<i>v6groupnetmask</i>	Displays the specific group information. When <i>v6groupnetmask</i> is not specified, the SSM Mapping status and all SSM Mapping entries on the VR are displayed.
vr <i>vr_name</i>	Specifies the virtual router name.

Default

N/A.

Usage Guidelines

Use this command to display the status of MLD-SSM mapping feature on a VR (if it is enabled or disabled), or to display the MLD-SSM mapping entries.

Example

The following command displays MLD SSM mapping on a VR :

```
# show mld ssm-map
MLD SSM mapping : Enabled
Group
-----
      Flags      Source
      -----
ff33::1/128
      d          2001:0DB8:1::3
      d          2001:0DB8:1::4
ff33::2/128
      -          2001:0DB8:1::5
      -          2001:0DB8:1::6

Flags : (d) Dynamic address obtained from DNS server
Total Entries : 6
```

When *v6groupnetmask* is specified, the SSM Mapping status and the SSM Mapping entries specific to the group range on the VR are displayed.

The following example displays a group containing only DNS name. “DNS” in parenthesis indicates the source is learned from the DNS server.

```
# show mld ssm-map ff33::1/128
MLD SSM mapping : Enabled

Group           : ff33::1/128
DNS Name        : abc
DNS Age         : 1512 seconds
Sources         : 2001:0DB8:1::3 (DNS)
                  2001:0DB8:1::4 (DNS)

Total Entries   : 4
```

The following example displays a group not configured with DNS name.

```
# show mld ssm-map ff32::1/128
MLD SSM mapping : Enabled
```

```
Group      : ff32::1/128
Sources    : 2001:0DB8:1::5
            : 2001:0DB8:1::6
Total Entries : 2
```

History

This command was first available in ExtremeXOS 15.5.

Platform Availability

This command is available on the platforms listed for the IPv6 multicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

show mpls

```
show mpls
```

Description

Displays *MPLS* configuration information.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

The `show mpls` command displays the current values for all the MPLS configuration parameters that apply to the entire switch. The parameters displayed include:

- MPLS and MPLS protocol (RSVP-TE and LDP) status.
- *SNMP* traps configuration.
- EXP examination/replacement configuration.
- the MPLS LSR ID.
- the list of the VLANs which have been added to MPLS.

Example

The following command displays the MPLS configuration parameters for the switch:

```
# show mpls
MPLS System
MPLS System      : User VR (Green)
MPLS Admin Status : Enabled
```



```
MPLS Oper Status      : Enabled
RSVP-TE Admin Status  : Enabled
RSVP-TE Oper Status   : Enabled
LDP Admin Status      : Enabled
LDP Oper Status       : Enabled
SNMP Traps            : Disabled
L2VPN SNMP Traps      : Enabled
EXP Examination       : Disabled
EXP Replacement       : Disabled
LSR ID                : 11.100.100.20
```

History

This command was first available in ExtremeXOS 11.6

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

show mpls bfd

```
show mpls bfd [{vlan} vlan_name | ip_addr]
```

Description

Displays MPLS BFD client information for a VLAN or interface.

Syntax Description

<i>vlan_name</i>	Specifies the VLAN for which to display MPLS BFD client information.
ip_addr	Specifies the IP address of an interface for which to display MPLS client information.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays the MPLS BFD client information for all next-hop peers:

```
Switch.1 # show mpls bfd
Next Hop IP      Count  Flags   Admin   Oper    IfName
-----
192.84.86.2      13   ASIU    Up      Up      vlan1
```

```
192.84.93.12      13   ASIU      Up      Up      vlan2
Flags: A=Session added to BFD server, S=BFD Server synced,
I=Session Init complete, U=State Updates accepted
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

show mpls exp examination

```
show mpls exp examination
```

Description

Displays MPLS EXP value to QoS profile mappings and whether MPLS EXP examination is enabled or disabled.

Syntax Description

This command has no arguments or keywords.

Default

N/A.

Usage Guidelines

This command displays MPLS EXP value to QoS profile mappings and the status of MPLS EXP examination (enabled or disabled). These values are set using the `configure mpls exp examination qosprofile` command and can be reset using the `unconfigure mpls exp examination` command.

Example

The following is an example of the output of the `show mpls exp examination` command:

```
# show mpls exp examination
EXP --> QoS Profile mapping:
00 --> QP1
01 --> QP2
02 --> QP3
03 --> QP4
04 --> QP5
05 --> QP6
06 --> QP7
```

```
07 --> QP8
EXP Examination is disabled
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

show mpls exp replacement

```
show mpls exp replacement
```

Description

Displays the *MPLS QoS* profile to EXP value mappings and the status of MPLS EXP replacement (enabled or disabled).

Syntax Description

This command has no arguments or keywords.

Default

N/A.

Usage Guidelines

This command displays MPLS QoS profile to EXP value mappings and the status of MPLS EXP replacement (enabled or disabled). These values are set using the `configure mpls exp replacement qosprofile` command and can be reset using the `unconfigure mpls exp replacement` command.

Example

The following is an example of the output of the `show mpls exp replacement` command:

```
* BD-10K.28 # show mpls exp replacement
QoS Profile --> EXP mapping:
QP1 --> 00
QP2 --> 01
QP3 --> 02
QP4 --> 03
QP5 --> 04
QP6 --> 05
QP7 --> 06
QP8 --> 07
EXP Replacement is disabled
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

show mpls interface

```
show mpls interface [{vlan} vlan_name] [detail]
```

Description

Displays the *MPLS* interface information. Information is displayed in tabular format for all *VLANs* that have been added to MPLS.

Syntax Description

vlan	Specifies to display information for one VLAN.
<i>vlan_name</i>	Specifies the name of the VLAN.
detail	Specifies to display additional status information about the interface.

Default

N/A.

Usage Guidelines

Information displayed includes:

- a list of all VLANs added to MPLS.
- MTU size.
- Local interface IP address.
- Number of RSVP-TE neighbors.
- Number of LDP adjacencies.
- RSVP-TE and LDP uptimes.
- MPLS protocols and capabilities configured on each VLAN.

Specifying the optional detail keyword displays the information in verbose form and also includes the operational state for RSVP-TE and LDP. Specifying a VLAN limits the output to that of the individual VLAN.

Example

The following command display MPLS interface information:

```
Switch # show mpls interface
Local          RSVP-TE      LDP
VLAN Name      IP Address    MTU    UpTm  #Nbr  UpTm  #Adj  Flags
-----
loopb          11.100.100.218 1500   3h    0     3h    0     MRL-I-U
toratora       192.84.86.1    1500   3h    0     3h    1     MRL-IBU
tordoze        192.84.93.1    1500   3h    0     3h    1     MRL-IbU
torfour        192.84.83.1    1500   --    0     --    0     MRL-I--
torinasp       192.84.85.1    1500   --    0     --    0     MRL-I--
Flags: (M) MPLS Enabled, (R) RSVP-TE Enabled, (L) LDP Enabled,
(P) PHP Enabled, (I) IP Forwarding Enabled, (B) BFD Enabled,
(b) BFD Disabled(Sessions Exist) (U) MPLS Operational
```

The following command displays detailed MPLS information for VLAN 1:

```
# show mpls interface vlan1 detail
VLAN Name : vlan1
Local IP Address      : 192.84.86.1
IP Forwarding         : Enabled
MPLS I/F MTU         : 1500
PHP Status            : Disabled
BFD Status            : Enabled
MPLS Admin Status     : Enabled
MPLS Oper Status      : Disabled
RSVP-TE Admin Status  : Enabled
Oper Status           : Disabled
UpTime                : 0d:0h:0m:0s
# Neighbors           : 0
LDP Admin Status      : Enabled
Oper Status           : Disabled
UpTime                : 0d:0h:0m:0s
# Link Adjacencies    : 0
```

History

This command was first available in ExtremeXOS 11.6.

The BFD flags and status were first available in ExtremeXOS 12.4.

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

show mpls label

```
show mpls {rsvp-te | static} label {summary | label_num | [advertised |
received] {label_num} | received implicit-null}
```

Description

Displays label information for all label types and protocols.

Syntax Description

rsvp-te	Specifies that only RSVP-TE LSP labels are displayed.
static	Specifies that only static LSP labels are displayed.
summary	Specifies that summary information for the labels is displayed.
<i>label_num</i>	Specifies that only labels that match the specified hexadecimal label number are displayed.
advertised	Specifies that only labels advertised to other routers are displayed.
received	Specifies that only labels received from other routers are displayed.
implicit-null	Specifies that only implicit-null labels are displayed.

Default

N/A.

Usage Guidelines

If no options are specified, tabular information for all labels, except advertised implicit-null labels, is displayed. The following table describes the display fields that appear when this command is entered.

Table 30: show mpls label Field Definitions

Field	Definition
Advertised Label	Advertised Label is the label advertised to other routers.
Destination Mapping	Destination Mapping displays the destination IP address and mask.
LSP Flags	LSP Flags lists the flags for each label, which are described in the key at the bottom of the display.
Peer Label	The peer label that replaces the received label of incoming packets during a label swap on a transit LSP. The peer label appears only for transit LSPs.
NHop Type	NHop Type helps define the handling of a packet arriving with the specified advertised label. Possible values are: IP—The label on the packet is swapped and the packet forwarded to the IP address shown in the NextHop field. locl—Indicates that the tunnel destination is local to this switch. VLAN—The label on the packet is stripped and is IP routed according to the Destination Mapping field. VRF—For advertised labels, the Next Hop column contains the name of the virtual router to which packets with the given Layer 3 VPN label will be forwarded. For received labels, the Next Hop column displays the router ID of the BGP peer, and the Name column displays the name of the VR using this label. VPLS—The label on the packet is stripped and forwarded according to the service configuration of the VPLS specified in the Name field.
Next Hop	Displays an interface ID for the next hop router, which can be a VLAN name, VMAN name, or IP address.

Table 30: show mpls label Field Definitions (continued)

Field	Definition
Name	Displays a VPLS or LSP name.
Received Label	Received Label displays the label received from other routers.

**Note**

Unsupported labels will contain no information.

Example

The following command displays all labels except received implicit-null labels:

```
* Switch.1 # show mpls label
```

Advertised Label	Destination Mapping	LSP Flags	Peer Label	NHOp Type	NextHop	Name

0x8082c	3.3.3.3/32	-LE	--	VLAN	lpbk	--
0x8082d	5.5.5.5/32	2LE	--	VPLS	--	extreme-501
0x8082e	5.5.5.5/32	2LE	--	VPLS	--	extreme-503
0x8082f	5.5.5.5/32	2LE	--	VPLS	--	extreme-504
0x80830	5.5.5.5/32	2LE	--	VPLS	--	extreme-505
0x80831	5.5.5.5/32	2LE	--	VPLS	--	extreme-506
0x80832	5.5.5.5/32	2LE	--	VPLS	--	extreme-507
0x80833	5.5.5.5/32	2LE	--	VPLS	--	extreme-508
0x80834	5.5.5.5/32	2LE	--	VPLS	--	extreme-509
0x8082a	3.3.3.3/32	-RE	--	loc1	3.3.3.3	lsp5to3-2
0x8082b	101.0.0.1/32	-RE	--	loc1	101.0.0.1	lsp5to3
0x80400	--	3-E	--	VRF	blue-vr	--
0x80401	--	3-E	--	VRF	red-vr	--
0x80402	--	3-E	--	VRF	white-vr	--

Received Label	Destination Mapping	LSP Flags	NextHop	Name

0x8082d	5.5.5.5/32	2LI	5.5.5.5	extreme-501
0x8082f	5.5.5.5/32	2LI	5.5.5.5	extreme-503
0x80830	5.5.5.5/32	2LI	5.5.5.5	extreme-504
0x80831	5.5.5.5/32	2LI	5.5.5.5	extreme-505
0x80832	5.5.5.5/32	2LI	5.5.5.5	extreme-506
0x80833	5.5.5.5/32	2LI	5.5.5.5	extreme-507
0x80834	5.5.5.5/32	2LI	5.5.5.5	extreme-508
0x80835	5.5.5.5/32	2LI	5.5.5.5	extreme-509
0x8082a	101.0.0.2/32	-RI	101.0.0.2	lsp3to5
0x8082b	5.5.5.5/32	-RI	101.0.0.2	lsp3to5-2
0x80401	--	3-I	5.5.5.5	red-vr
0x80400	--	3-I	5.5.5.5	red-vr
0x80401	--	3-I	5.5.5.5	blue-vr
0x80400	--	3-I	5.5.5.5	blue-vr
0x80402	--	3-I	5.5.5.5	white-vr

Flags: (3) L3VPN, (2) L2VPN, (L) LDP, (R) RSVP-TE, (S) Static
 (T) Transit LSP, (I) Ingress to LSP, (E) Egress from LSP,
 (M) Multiple Next Hops

Summary of Labels	Advertised	Received
Total number of RSVP-TE LSP labels	2	2

Total number of LDP LSP labels	1	0
Total number of Static LSP labels	0	0
Total number of L2VPN Labels	8	8
Total number of L3VPN Labels	3	5

The following command displays all rsvp-te labels except received implicit-null labels:

```
* Switch.2 # show mpls rsvp-te label
```

Advertised Label	Destination Mapping	Label Flags	Peer Label	NHop Type	NextHop	Name
0x80834	101.0.0.1/32	-RE	--	loc1	101.0.0.1	lsp5to3
0x80835	3.3.3.3/32	-RE	--	loc1	3.3.3.3	lsp5to3-2

Received Label	Destination Mapping	Label Flags	NextHop	Name
0x8082a	101.0.0.2/32	-RI	101.0.0.2	lsp3to5
0x8082b	5.5.5.5/32	-RI	101.0.0.2	lsp3to5-2

Flags: (3) L3VPN, (2) L2VPN, (L) LDP, (R) RSVP-TE, (S) Static, (T) Transit LSP, (I) Ingress to LSP, (E) Egress from LSP, (M) Multiple Next Hops

Summary of Labels	Advertised	Received
Total number of RSVP-TE LSP labels	2	2

History

This command was first available in ExtremeXOS 11.6.

This command was modified to display only RSVP-TE and static label information in ExtremeXOS 12.5. Additional commands were added to display LDP Layer 2 VPN and Layer 3 VPN labels.

Platform Availability

This command is available only on the platforms that support *MPLS* as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

show mpls label l3vpn

```
show mpls label l3vpn {summary | label_num | [advertised | received]
                       {label_num}}
```

Description

Displays Layer 3 VPN label information.

Syntax Description

summary	Specifies that summary information for the labels is displayed.
<i>label_num</i>	Specifies that only labels that match the specified hexadecimal label number are displayed.
advertised	Specifies that only labels advertised to other routers are displayed.
received	Specifies that only labels received from other routers are displayed.

Default

N/A.

Usage Guidelines

If no options are specified, tabular information for all Layer 3 VPN labels is displayed. The following table describes the display fields that appear when this command is entered.

Table 31: show mpls label l3vpn Field Definitions

Field	Definition
Advertised Label	Advertised Label is the label advertised to other routers.
Destination Mapping	Destination Mapping displays the destination IP address and mask.
LSP Flags	LSP Flags lists the flags for each label, which are described in the key at the bottom of the display.
Peer Label	The peer label that replaces the received label of incoming packets during a label swap on a transit LSP. The peer label appears only for transit LSPs.
NHop Type	NHop Type helps define the handling of a packet arriving with the specified advertised label. Possible values are: VRF—For advertised labels, the Next Hop column contains the name of the virtual router to which packets with the given Layer 3 VPN label will be forwarded. For received labels, the Next Hop column displays the router ID of the <i>BGP</i> peer, and the Name column displays the name of the VR using this label.
Next Hop	Displays an interface ID for the next hop router, which can be a <i>VLAN</i> name, vMAN name, or IP address.
Name	Displays the name of the VR using this label.
Received Label	Received Label displays the label received from other routers.

Example

The following command displays all Layer 3 VPN labels:

```
* Switch.1 # show mpls label l3vpn
```

Advertised Label	Destination Mapping	LSP Flags	Peer Label	NHop Type	NextHop	Name
0x80400	--	3-E	--	VRF	blue-vr	--

```

0x80401      --      3-E      --  VRF    red-vr      --
0x80402      --      3-E      --  VRF    white-vr     --
Received      Destination  LSP
Label          Mapping  Flags    NextHop      Name
-----
0x80400      --      3-I    3.3.3.3      red-vr
0x80400      --      3-I    3.3.3.3      blue-vr
0x80401      --      3-I    3.3.3.3      red-vr
0x80401      --      3-I    3.3.3.3      blue-vr
0x80402      --      3-I    3.3.3.3      white-vr
Flags: (3) L3VPN, (2) L2VPN, (L) LDP, (R) RSVP-TE, (S) Static, (P) LSP
(T) Transit LSP, (I) Ingress to LSP, (E) Egress from LSP,
(M) Multiple Next Hops
Total number of L3VPN Labels advertised      : 3
Total number of L3VPN Labels received        : 5

```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

show mpls label usage

```
show mpls label usage {static}
```

Description

Displays the label ranges on the current running system, including configurable and non-configurable ranges.

Syntax Description

usage	Displays MPLS label ranges on the current running system, including configurable and non-configurable ranges.
static	Displays MPLS label usage of static labels.

Default

N/A.

Usage Guidelines

With the addition of the static PW configuration, there is the need to configure static labels and display more detailed label information. This command displays the label ranges on the current running system, including configurable and non-configurable ranges. The output also includes hardware resource usage to provide better information about MPLS hardware utilization and capacity.

Example

The following command displays label ranges in-use, information about the static label usage, and some of the label hardware usage:

```
* show mpls lab usage
```

Label Type	Size	Label Range
Supported	1048576	0x00000 - 0xfffff (0 - 1048575)
Reserved	16	0x00000 - 0x0000f (0 - 15)
Static	300	0x00010 - 0x0013b (16 - 315)
L3VPN	255	0x0013c - 0x0023a (316 - 570)
Dynamic	7365	0x0023b - 0x01eff (571 - 7935)
Internal Use	256	0x01f00 - 0x01fff (7936 - 8191)


```
Static Label Configuration Usage
```

	In-Use	Avail	Total	%Avail
Total	24	276	300	92%
Ingress LSP	2			
Egress LSP	0			
Transit LSP	2			
L2VPN	17			
CES	3			


```
Label Hardware Resource Usage
```

	Incoming				Outgoing			
	In-Use	Avail	Total	%Avail	In-Use	Avail	Total	%Avail
Static Ingress LSP	-	-	-	-	2	0	0	0%
Static Transit LSP	0	278	300	92%	0	0	0	0%
Static Egress LSP	2	278	300	92%	-	-	-	-
Static L2VPN	17	278	300	92%	0	0	0	0%
RSVP-TE Ingress LSP	-	-	-	-	5	0	0	0%
RSVP-TE Transit LSP	0	7342	7365	99%	0	0	0	0%
RSVP-TE Egress LSP	1	7342	7365	99%	-	-	-	-
LDP Ingress LSP	-	-	-	-	5	0	0	0%
LDP Transit LSP	5	7342	7365	99%	5	0	0	0%
LDP Egress LSP	1	7342	7365	99%	-	-	-	-
LDP L2VPN	14	7342	7365	99%	14	0	0	0%
L3VPN	0	255	255	100%	0	0	0	0%

Here is an example of the **static** option:

```
Summit1.2 # show mpls label usage static
```

Inbound Label	Name	Peer/Dest	Flags
0x22 (34)	green01	11.100.100.241	2
0x23 (35)	green01	11.100.100.243	2
0x24 (36)	green01	11.100.100.244	2
0x27 (39)	green03	11.100.100.241	2
0x28 (40)	green03	11.100.100.243	2
0x31 (49)	red02	11.100.100.101	2
0x32 (50)	red02	11.100.100.102	2
0x34 (52)	red02	11.100.100.241	2
0x35 (53)	red02	11.100.100.242	2
0x36 (54)	red02	11.100.100.243	2
0x37 (55)	red02	11.100.100.244	2
0x3d (61)	red05	11.100.100.241	2

```

0x3f (63)      red05      11.100.100.244  2
0x81 (129)     red04      11.100.100.101  2
0x82 (130)     red04      11.100.100.102  2
0x84 (132)     red04      11.100.100.241  2
0x85 (133)     red04      11.100.100.242  2
0x86 (134)     red04      11.100.100.243  2
0x87 (135)     red04      11.100.100.244  2
0xd1 (209)     red05      11.100.100.101  2
0x104 (260)    lspstat02  11.100.100.101  L
0x10d (269)    lspstat04  11.100.100.102  L

Flags: (2) L2VPN, (C) CES, (L) Static LSP

Static Label Range:      0x10 - 0x13b (16 - 315)
Static Label Statistics: 300 total, 22 in-use, 278 (92%) available

```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

show mpls ldp

```
show mpls ldp
```

Description

Displays summary configuration and status information for LDP. Global status of LDP, LDP session timer configuration, loop detection, and label advertisement status are included in the display output.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command displays the global status of LDP, LDP session timer configuration, loop detection, label advertisement, and LDP-enabled VLANs.

The following table describes the display fields that appear when this command is entered.

Table 32: show mpls ldp Field Definitions

Field	Definition
LDP Admin Status	LDP Admin Status shows whether LDP has been administratively Enabled or Disabled .
LDP Oper Status	LDP Oper Status shows whether LDP is operating (Enabled) or not (Disabled).
Protocol Version	Protocol Version specifies the LDP protocol version number.
Label Retention Mode	Label Retention Mode is either Conservative or Liberal , as described in RFC 3036, LDP Specification. The ExtremeXOS software only supports Liberal mode. In Liberal mode, a label switch router maintains all received label-to-FEC mapping advertisements.
Label Distribution Method	Label Distribution Method is either Downstream Unsolicited or Downstream On Demand , as described in RFC 3036, LDP Specification. The ExtremeXOS software only supports Downstream Unsolicited mode. In Downstream Unsolicited mode, a label switch router distributes label bindings to peer label switch routers without waiting for the bindings to be requested.
Label Distribution Control Mode	Label Distribution Control Mode is either Independent or Ordered , as described in RFC 3036, LDP Specification. The ExtremeXOS software only supports Ordered mode. In Ordered control mode, a label switch router only creates a binding of a label-to-FEC when it is the egress router or it has received a label binding for that FEC from the next hop router for that FEC.
LDP <i>BGP</i> LSPs	LDP BGP LSPs shows whether LDP uses BGP routes. When the displayed value is Enabled , LDP accepts BGP routes and stores them in the LDP internal routing table. When the displayed value is Disabled , LDP does not accept BGP routes, which reduces memory requirements when LSPs based on BGP routes are not desired. Note that when Disabled is displayed, no LDP LSPs are established to prefixes for which BGP is the preferred routing protocol.
LDP Loop Detection	LDP Loop Detection displays the LDP loop detection configuration.
LDP Targeted Timers	LDP Targeted Timers displays the LDP timer configuration used for LDP targeted adjacencies and sessions.
LDP Link Timers	LDP Link Timers displays the LDP timer configuration used for LDP link adjacencies and sessions.

Example

The following command displays summary configuration and status information for LDP:

```
# show mpls ldp
LDP Admin Status           : Enabled
LDP Oper Status            : Enabled
Protocol Version           : v1*
Label Retention Mode       : Liberal*
Label Distribution Method   : Downstream Unsolicited*
Label Distribution Control Mode : Ordered*
LDP BGP LSPs              : Enabled
```

```

LDP Loop Detection
Status          : Disabled
Hop-Count Limit : 255
Path-Vector Limit : 255
LDP Targeted Timers
Hello Hold      : 45 seconds
Keep Alive Hold : 60 seconds
LDP Link Timers
Hello Hold      : 15 seconds
Keep Alive Hold : 40 seconds
Label Advertisement
Direct : Matching LSR-ID Only
Rip    : None
Static : None
LDP VLANs : vlan1
: vlan2
: vlan3
* Indicates parameters that cannot be modified
# show mpls ldp
LDP Admin Status          : Enabled
LDP Oper Status           : Disabled
Protocol Version          : vl*
Label Retention Mode      : Liberal*
Label Distribution Method  : Downstream Unsolicited*
Label Distribution Control Mode : Ordered*
LDP BGP LSPs              : Disabled
LDP Loop Detection
Status          : Disabled
Hop-Count Limit : 255
Path-Vector Limit : 255
LDP Targeted Timers
Hello Hold      : 45 seconds
Keep Alive Hold : 60 seconds
LDP Link Timers
Hello Hold      : 15 seconds
Keep Alive Hold : 40 seconds
Label Advertisement
Direct : Matching LSR-ID Only
Rip    : None
Static : None
LDP VLANs : karen
: lb
* Indicates parameters that cannot be modified

```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

show mpls ldp interface

```
show mpls ldp interface {{vlan} vlan_name} {detail | counters}
```

Description

Displays LDP information about *MPLS* interfaces. Summary information is displayed in tabular format for all *VLANs* that are configured for MPLS.

Syntax Description

vlan	Displays LDP interface information for one VLAN.
<i>vlan_name</i>	Specifies the name of the VLAN.
detail	Displays LDP interface information including LDP control packet counts.
counters	Displays only the LDP control protocol packet counts.

Default

N/A.

Usage Guidelines

If the optional **detail** keyword is specified, the information is shown in verbose form and LDP control packet counts are displayed. If the optional **counters** keyword is specified, only the LDP control protocol packet counts are shown. The counters are described in RFC 3036, LDP Specification.

Example

The following command displays detailed LDP information for the interface associated with VLAN 1:

```
# show mpls ldp interface vlan1 detail
VLAN Name : vlan1
Local IP Address      : 11.121.96.20
MPLS Admin Status    : Enabled
MPLS Oper Status     : Enabled
LDP Admin Status     : Enabled
LDP Oper Status      : Enabled
LDP UpTime           : 0d:1h:59m:56s
Current Adjacencies   : 1
Negotiated Hello Hold Time : 15000 ms
Time to Send Next Hello : 4060 ms
Link      Targeted
Counter          Adjacencies  Adjacencies
-----
Shutdown Notifications (Rcvd)      0          0
Shutdown Notifications (Sent)      0          0
Failed Session Attempts (NAKs)     0          0
Hello Errors                       0          0
Parameters Advertised Errors       0          0
Max PDU Length Errors              0          0
Label Range Errors                 0          0
Bad LDP ID Errors                  0          0
Bad PDU Length Errors              0          0
Bad Msg Length Errors              0          0
Bad TLV Length Errors              0          0
Bad TLV Value Errors               0          0
Keep-Alive Timeout Errors          0          0
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

show mpls ldp label

```
show mpls {ldp} label {lsp} {summary | label_num | [advertised |  
received] {label_num} | received implicit-null}
```

Description

Displays LDP LSP label information.

Syntax Description

ldp	Specifies that only LDP labels are displayed.
summary	Specifies that summary information for the labels is displayed.
<i>label_num</i>	Specifies that only labels that match the specified hexadecimal label number are displayed.
advertised	Specifies that only labels advertised to other routers are displayed.
received	Specifies that only labels received from other routers are displayed.
implicit-null	Specifies that only implicit-null labels are displayed.

Default

N/A.

Usage Guidelines

If no options are specified, tabular information for all LDP labels, except advertised implicit-null labels, is displayed. The following table describes the display fields that appear when this command is entered.

Table 33: show mpls ldp label Field Definitions

Field	Definition
Advertised Label	Advertised Label is the label advertised to other routers.
Destination Mapping	Destination Mapping displays the destination IP address and mask.
LSP Flags	LSP Flags lists the flags for each label, which are described in the key at the bottom of the display.
Peer Label	The peer label that replaces the received label of incoming packets during a label swap on a transit LSP. The peer label appears only for transit LSPs.

Table 33: show mpls ldp label Field Definitions (continued)

Field	Definition
NHop Type	NHop Type helps define the handling of a packet arriving with the specified advertised label. Possible values are: IP—The label on the packet is swapped and the packet forwarded to the IP address shown in the NextHop field.vlan—The label on the packet is stripped and is IP routed according to the Destination Mapping field.local—Indicates that the tunnel destination is local to this switch.
Next Hop	Displays an interface ID for the next hop router, which can be a VLAN name, VMAN name, or IP address.
Name	Displays an LSP name.
Received Label	Received Label displays the label received from other routers.

Example

The following command displays all LDP labels except received implicit-null labels:

```
* Switch.1 # show mpls ldp label
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

show mpls ldp label advertised

```
show mpls ldp label {lsp} advertised implicit-null {ipNetmask}
```

Description

Displays advertised LDP LSP implicit-null label information.

Syntax Description

<i>ipNetmask</i>	Specifies an IP network mask.
------------------	-------------------------------

Default

N/A.

Usage Guidelines

If no options are specified, tabular information for all LDP advertised implicit-null labels is displayed. The following table describes the display fields that appear when this command is entered.

Table 34: show mpls ldp label Field Definitions

Field	Definition
Advertised Label	Advertised Label is the label advertised to other routers.
Destination Mapping	Destination Mapping displays the destination IP address and mask.
LSP Flags	LSP Flags lists the flags for each label, which are described in the key at the bottom of the display.
Peer Label	The peer label that replaces the received label of incoming packets during a label swap on a transit LSP. The peer label appears only for transit LSPs.
NHop Type	NHop Type helps define the handling of a packet arriving with the specified advertised label. Possible values are: IP—The label on the packet is swapped and the packet forwarded to the IP address shown in the NextHop field. vlan—The label on the packet is stripped and is IP routed according to the Destination Mapping field. locl—Indicates that the tunnel destination is local to this switch.
Next Hop	Displays an interface ID for the next hop router, which can be a VLAN name, VMAN name, or IP address.
Name	Displays a VPLS or LSP name.
Received Label	Received Label displays the label received from other routers.

Example

The following command displays all advertised LDP implicit-null labels:

```
* Switch.1 # show mpls ldp label advertised implicit-null
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in in the [Switch Engine 32.3 Feature License Requirements](#) document.

show mpls ldp label l2vpn retained

```
show mpls ldp label l2vpn retained {ipaddress}
```

Description

Displays Layer 2 VPN liberally retained labels received from a peer.

Syntax Description

<i>ipaddress</i>	Specifies an IP address.
------------------	--------------------------

Default

N/A.

Usage Guidelines

If no options are specified, tabular information for all Layer 2 VPN liberally-retained labels is displayed. The following table describes the display fields that appear when this command is entered.

Table 35: show mpls ldp label l2vpn retained Field Definitions

Field	Definition
Advertised Label	Advertised Label is the label advertised to other routers.
Destination Mapping	Destination Mapping displays the destination IP address and mask.
LSP Flags	LSP Flags lists the flags for each label, which are described in the key at the bottom of the display.
Peer Label	The peer label that replaces the received label of incoming packets during a label swap on a transit LSP. The peer label appears only for transit LSPs.
NHop Type	NHop Type helps define the handling of a packet arriving with the specified advertised label. Possible values are: VPLS—The label on the packet is stripped and forwarded according to the service configuration of the VPLS specified in the Name field.
Next Hop	Displays an interface ID for the next hop router, which can be a VLAN name, VMAN name, or IP address.
Name	Displays a VPLS name.
Received Label	Received Label displays the label received from other routers.

Example

The following command displays liberally-retained Layer 2 VPN labels received from peers:

```
show mpls ldp label l2vpn retained
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

show mpls ldp label l2vpn

```
show mpls {ldp} label l2vpn {summary | label_num | [advertised |  
received] {label_num}}
```

Description

Displays LDP Layer 2 VPN label information.

Syntax Description

ldp	Specifies that only LDP labels are displayed.
summary	Specifies that summary information for the labels is displayed.
<i>label_num</i>	Specifies that only labels that match the specified hexadecimal label number are displayed.
advertised	Specifies that only labels advertised to other routers are displayed.
received	Specifies that only labels received from other routers are displayed.

Default

N/A.

Usage Guidelines

If no options are specified, tabular information for all LDP Layer 2 VPN labels is displayed. The following table describes the display fields that appear when this command is entered.

Table 36: show mpls ldp label l2vpn Field Definitions

Field	Definition
Advertised Label	Advertised Label is the label advertised to other routers.
Destination Mapping	Destination Mapping displays the destination IP address and mask.
LSP Flags	LSP Flags lists the flags for each label, which are described in the key at the bottom of the display.
Peer Label	The peer label that replaces the received label of incoming packets during a label swap on a transit LSP. The peer label appears only for transit LSPs.
NHop Type	NHop Type helps define the handling of a packet arriving with the specified advertised label. Possible values are: VPLS—The label on the packet is stripped and forwarded according to the service configuration of the VPLS specified in the Name field.
Next Hop	Displays an interface ID for the next hop router, which can be a <u>VLAN</u> name, VMAN name, or IP address.
Name	Displays a VPLS or LSP name.
Received Label	Received Label displays the label received from other routers.

Example

The following command displays all Layer 2 VPN labels:

```
* Switch.1 # show mpls ldp label l2vpn
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available only on the platforms that support *MPLS* as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

show mpls ldp label lsp retained

```
show mpls ldp label lsp retained {ipNetmask}
```

Description

Displays LSP liberally retained labels received from a peer.

Syntax Description

<i>ipNetmask</i>	Specifies an IP network mask.
------------------	-------------------------------

Default

N/A.

Usage Guidelines

If no options are specified, tabular information for all LSP liberally-retained labels is displayed. The following table describes the display fields that appear when this command is entered.

Table 37: show mpls ldp label lsp retained Field Definitions

Field	Definition
Advertised Label	Advertised Label is the label advertised to other routers.
Destination Mapping	Destination Mapping displays the destination IP address and mask.
LSP Flags	LSP Flags lists the flags for each label, which are described in the key at the bottom of the display.
Peer Label	The peer label that replaces the received label of incoming packets during a label swap on a transit LSP. The peer label appears only for transit LSPs.

Table 37: show mpls ldp label lsp retained Field Definitions (continued)

Field	Definition
NHop Type	NHop Type helps define the handling of a packet arriving with the specified advertised label. Possible values are: IP—The label on the packet is swapped and the packet forwarded to the IP address shown in the NextHop field.vlan—The label on the packet is stripped and is IP routed according to the Destination Mapping field.local—Indicates that the tunnel destination is local to this switch.
Next Hop	Displays an interface ID for the next hop router, which can be a VLAN name, VMAN name, or IP address.
Name	Displays an LSP name.
Received Label	Received Label displays the label received from other routers.

Example

The following command displays liberally-retained LSP labels received from peers:

```
* Switch.1 # show mpls ldp label lsp retained lsp
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

show mpls ldp label retained

```
show mpls ldp label retained [l2vpn {ipaddress} | lsp {ipNetmask}]
```

Description

Displays liberally-retained labels received from a peer for either the Layer 2 VPN protocol or LSP protocol.

Syntax Description

<i>ipaddress</i>	Specifies an IP address.
<i>ipNetmask</i>	Specifies an IP network mask.

Default

N/A.

Usage Guidelines

If no options are specified, tabular information for all liberally-retained labels is displayed. The following table describes the display fields that appear when this command is entered.

Table 38: show mpls label retained Field Definitions

Field	Definition
Advertised Label	Advertised Label is the label advertised to other routers.
Destination Mapping	Destination Mapping displays the destination IP address and mask.
LSP Flags	LSP Flags lists the flags for each label, which are described in the key at the bottom of the display.
Peer Label	The peer label that replaces the received label of incoming packets during a label swap on a transit LSP. The peer label appears only for transit LSPs.
NHOp Type	NHOp Type helps define the handling of a packet arriving with the specified advertised label. Possible values are: IP—The label on the packet is swapped and the packet forwarded to the IP address shown in the NextHop field. local—Indicates that the tunnel destination is local to this switch. VLAN—The label on the packet is stripped and is IP routed according to the Destination Mapping field. VPLS—The label on the packet is stripped and forwarded according to the service configuration of the VPLS specified in the Name field.
Next Hop	Displays an interface ID for the next hop router, which can be a VLAN name, VMAN name, or IP address.
Name	Displays a VPLS or LSP name.
Received Label	Received Label displays the label received from other routers.

Example

The following command displays liberally-retained LSP labels received from peers:

```
# show mpls ldp label retained lsp
```

Prefix	Peer LSR ID	Peer Label	Next Hop	VLAN
30.30.30.0/24	3.3.3.3	0x00435	30.30.30.2	r1-r3
2.2.2.2/32	3.3.3.3	0x00439	30.30.30.2	r1-r3
22.22.22.22/32	3.3.3.3	0x0043a	30.30.30.2	r1-r3
10.10.10.0/24	22.22.22.22	0x00434	10.10.10.2	r1-r2
30.30.30.0/24	22.22.22.22	0x00438	10.10.10.2	r1-r2
1.1.1.1/32	22.22.22.22	0x00439	10.10.10.2	r1-r2
3.3.3.3/32	22.22.22.22	0x0043a	10.10.10.2	r1-r2

History

This command was first available in ExtremeXOS 12.5.

Peer LSR information added in ExtremeXOS 22.5.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

show mpls ldp lsp

```
show mpls ldp lsp {prefix ipNetmask} [ingress | egress | transit] [detail]
```

Description

Displays the LSP information associated with LDP that is used to forward packets within the [MPLS](#) network. If no options are specified, summary information for all LSPs is displayed.

Syntax Description

prefix	Displays information for a single FEC that matches the prefix.
<i>ipNetmask</i>	Designates the FEC for which to display information.
ingress	Displays information for LSPs that originate from the switch into the MPLS network.
egress	Displays information for LSPs that terminate at the switch from the MPLS network.
transit	Displays information for LSPs that traverse the switch.
detail	Display detailed LSP information.

Default

N/A.

Usage Guidelines

If no options are specified, this command displays summary information for all LSPs.

Optionally, the LSPs displayed can be further qualified by the keywords **ingress**, **egress**, and **transit**. These keywords qualify the LSPs displayed from the perspective of the switch. Ingress LSPs originate from the switch into the MPLS network. Egress LSPs terminate at the switch from the MPLS network. Transit LSPs traverse the switch. If the optional **prefix** keyword is specified, only the LSP information associated with the FEC that matches the prefix is displayed.

If the **detail** keyword is specified, information is displayed in verbose form and includes received packet and byte counts.

Example

The following command displays LDP information for an ingress LSP:

```
# show mpls ldp lsp 11.100.100.59/32 ingress detail
```



```

FEC IP/Prefix: 11.100.100.59/32
Next Hop I/F      : m5vlan1
Next Hop Addr     : 12.224.0.55
Advertised Label  : n/a                Received Label : 0x80403 (525315)
Rx Packets        : n/a                Tx Packets      : 61
Rx Bytes          : n/a                Tx Bytes       : 4294967296

```

The following command displays LDP information for a transit LSP:

```

# show mpls ldp lsp 11.100.100.55/32 transit detail
FEC IP/Prefix: 11.100.100.55/32
Next Hop I/F      : m5vlan1
Next Hop Addr     : 12.224.0.55
Advertised Label  : 0x11 (17)          Received Label : 0x80403 (525315)
Rx Packets        : 61                  Tx Packets      : 61
Rx Bytes          : 4294967296          Tx Bytes       : 4294967296

```

The following command displays LDP information for an egress LSP:

```

# show mpls ldp lsp 11.100.100.30/32 egress detail
FEC IP/Prefix: 11.100.100.30/32
Direct VLAN       : loop
Advertised Label  : 0x80400 (525312)   Received Label : n/a
Rx Packets        : 61                  Tx Packets      : n/a
Rx Bytes          : 4294967296          Tx Bytes       : n/a

```

History

This command was first available in ExtremeXOS 11.6.

The output for this command was modified in ExtremeXOS 12.2.2.

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

show mpls ldp peer

```
show mpls ldp peer {ipaddress} {detail}
```

Description

Displays information about the status of the LDP sessions and hello adjacencies for all LDP peers.

Syntax Description

<i>ipaddress</i>	Display session and hello adjacency information for a single LDP peer.
detail	Display additional detailed information related to the session and adjacencies.

Default

N/A.

Usage Guidelines

Specifying the LDP peer's ipaddress displays session and hello adjacency information for a single LDP peer. When the `detail` keyword is specified, additional detailed information related to the session and adjacencies is displayed.

[Table 30](#) on page 2890 describes the display fields that appear when this command is entered.

Table 39: show mpls ldp peer Field Definitions

Field	Definition
IP Address	Local IP address, which is used as the LSR-ID.
LDP Peer	LDP identifier of LDP peer.
State	Displays the state of the Initialization State Machine as described in RFC 3036, LDP Specification. The states are: NonExistent, Initialized, OpenRec, OpenSent, and Operational.
Uptime	Displays the total time the session has been operational.
Adjacencies	Displays the number of active adjacencies with the LDP peer.
Index	The Entity Index used in the LDP Entity Table MIB.
Targeted Peer	The IP address of the peer used in Extended Discovery. If this is not a targeted peer (Basic Discovery was used), this field displays Not Targeted.
Attempted Sessions	This and other counters are described in RFC 3036, LDP Specification.
Shutdown Notifications	This and other counters are described in RFC 3036, LDP Specification.
Peer	IP address of peer.
Peer Label Space	The label space as given by the peer and derived from the LDP Identifier. A zero value represents the global or per-platform label space. A non-zero value represents a per interface label space.
Session State	Displays the state of the Initialization State Machine as described in RFC 3036, LDP Specification. The states are: NonExistent, Initialized, OpenRec, OpenSent, and Operational.
Session Uptime	Displays the total time the session has been operational.
Discontinuity Time	The system uptime for the most recent period after which one or more of the session's counters suffered a discontinuity.
Keep Alive Hold Timer	Displays the configured keep alive hold timer value and the remaining hold time.
Label Distribution Method	Label Distribution Method is either DU (Downstream Unsolicited) or DOD (Downstream On Demand), as described in RFC 3036, LDP Specification.
Max PDU Length	The maximum allowable length for LDP PDUs (Protocol Data Units) for this session.
Unknown Msg Type Errors	The number of messages received for this session without a recognized message type and without the ignore bit set.

Table 39: show mpls ldp peer Field Definitions (continued)

Field	Definition
Unknown TLV Errors	The number of TLVs received for this session without a recognized TLV type and without the ignore bit set.
Next Hop Addr(s)	A list of next hop addresses received from the peer through LDP address messages.

Example

The following command displays *MPLS* LDP session information for the LDP entity 11.100.100.30:

```

Mariner3.59 # show mpls ldp peer
IP Address      LDP Peer      State      Uptime      Adjacencies
11.100.100.30   11.100.100.55:0 Operational  0d:14h:51m:53s  1
11.100.100.30   14.4.0.99:15  Operational  0d:1h:0m:43s   1
11.100.100.30   14.4.0.99:16  Operational  0d:0h:34m:51s   1
Adjacencies:
Index           : 1                      Attempted Sessions      : 1
Targeted Peer   : 11.100.100.210:0      Shutdown Notifications   : Sent 0
Rcvd 0

Mariner3.32 # show mpls ldp peer detail
Peer: 11.100.100.55      Peer label space: 0 (global)
Session State           : Operational
Session Uptime          : 0d:0h:13m:41s
...
Peer: 14.4.0.99          Peer label space: 15
Session State           : Operational
Session Uptime          : 0d:0h:57m:4s
...
Peer: 14.4.0.99          Peer label space: 16
Session State           : Operational
Session Uptime          : 0d:0h:31m:12s
...

* DUT65.2 # show mpls ldp peer detail
Peer: 11.100.100.210      Peer Label Space: 0 (global)
Session State           : Operational
Session Uptime          : 0d:0h:6m:30s
Discontinuity Time      : 34677
Keep Alive Hold Timer    : 40 (remaining: 37.86)
Label Distribution Method : DU
Max PDU Length          : 4096
Unknown Msg Type Errors  : 0
Unknown TLV Errors       : 0
Next Hop Addr(s)        : 11.100.100.210 12.20.20.210
Adjacencies:
Index           : 1                      Attempted Sessions      : 0
Targeted Peer   : 11.100.100.210:0      Shutdown Notifications   : Sent 0
Rcvd 0
Admin Status    : Enabled                No Hello Errors         : 0
Operational Status : Up                  Advertisement Errors    : 0
Label Retention Mode : Liberal                Max PDU Errors          : 0
Hop Count Limit    : Disabled              Bad LDP Identifier Errors : 0
Path Vector Limit   : Disabled              Bad PDU Length Errors   : 0
Hello Hold Timer    : 45 (remaining: 38)    Bad TLV Length Errors   : 0
Malformed TLV Errors : 0
Bad Message Length Errors : 0
Session Rejected Errors : 0
Keep Alive Expired Errors : 0

```

```

Index                : 6
Targeted Peer        : Not Targeted
Rcvd 0
Admin Status         : Enabled
Operational Status   : Up
Label Retention Mode : Liberal
Hop Count Limit      : Disabled
Path Vector Limit    : Disabled
Hello Hold Timer     : 15 (remaining: 10)
VLAN                  : v1
Interface address    : 12.20.20.182
Session Rejected Errors : 0
Keep Alive Expired Errors : 0
Attempted Sessions   : 0
Shutdown Notifications : Sent 0
No Hello Errors      : 0
Advertisement Errors  : 0
Max PDU Errors       : 0
Bad LDP Identifier Errors : 0
Bad PDU Length Errors : 0
Bad TLV Length Errors : 0
Malformed TLV Errors : 0
Bad Message Length Errors : 0

```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

show mpls rsvp-te bandwidth

```
show mpls rsvp-te bandwidth {{vlan} vlan_name} {detail}
```

Description

Displays the reserved bandwidth for each TE LSP by interface.

Syntax Description

<i>vlan_name</i>	Displays the reserved bandwidth for each TE LSP associated with the specified <u>VLAN</u> .
detail	Displays the path information in verbose format.

Default

All TE LSPs for all RSVP-TE-enabled interfaces are shown.

Usage Guidelines

This command displays the reserved bandwidth for each TE LSP by interface. By default, all TE LSPs for all RSVP-TE enabled interfaces are shown.



Note

Beginning with ExtremeXOS Release 12.2.1, the receive bandwidth can only be used for tracking. If the configured receive bandwidth is exceeded, the available bandwidth shown might be negative. In this case, “Ovr” is displayed to indicate that the link is oversubscribed in the receive direction. The **detail** option can be used to show the actual LSPs using this bandwidth.

The optional **vlan** keyword limits the display to only those LSPs that have bandwidth reservations against the specified VLAN. Only committed-rate bandwidth is displayed. Bandwidth is displayed as either received or transmitted bandwidth with respect to the switch.

LSPs are listed using the configured or signaled LSP name. If the LSP name was not included in the setup control messages (which can only occur when using OEM vendor equipment), the LSP is uniquely identified using a concatenated string that includes the tunnel ID and source IP address. Per VLAN, each LSP is listed in descending priority order. That is, the LSPs listed at the top of each VLAN have the highest bandwidth priority and are less likely to be preempted. Bandwidth priority is determined by the signaled hold-priority and the uptime. The TE LSP with a hold-priority of zero and the highest uptime has the highest bandwidth priority and the TE LSP with a hold-priority of seven and the lowest uptime has the lowest bandwidth priority.

Use the **detail** keyword to display detailed information.

Example

The following command displays bandwidth reservation information for the specified VLAN:

```
show mpls rsvp-te bandwidth vlan vlan_1 detail
# show mpls rsvp-te bandwidth vlan_1 detail
Vlan      Dir  Pool      CIR (per priority level)
LSP        0    1    2    3    4    5    6    7
-----
vlan_1      Rx   300
vlalsp1     -    -    10    -    -    -    -    -
-----
Available  300  300  290  290  290  290  290  290
Tx   500
vlalsp2     -    -    9    -    -    -    -    -
-----
Available  500  500  491  491  491  491  491  491
(Rx)Receive Bandwidth (Tx)Transmit Bandwidth
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

show mpls rsvp-te interface

```
show mpls rsvp-te interface {{vlan} vlan_name} {detail | counters}
```

Description

Displays RSVP-TE information about [MPLS](#) interfaces.

Syntax Description

vlan	Display information for one VLAN .
<i>vlan_name</i>	Specifies the name of the VLAN.
detail	Display RSVP-TE information including the interface up time and LDP control packet counts.
counters	Display only the RSVP-TE control protocol packet counts.

Default

N/A.

Usage Guidelines

This command displays RSVP-TE information about MPLS interfaces. Summary information is displayed in tabular format for all VLANs that are configured for MPLS. The following information is displayed:

- VLAN name.
- Bandwidth reserved.
- TE metric.
- Hello interval time.
- Refresh interval time.
- Summary refresh time.
- Bundle message time.
- Uptime.
- Number of neighbors.
- RSVP-TE state information

When the optional **detail** keyword is specified, additional RSVP-TE information is displayed. This additional information includes:

- RSVP-TE hello keep multiplier.
- RSVP-TE refresh keep multiplier.

- RSVP-TE available bandwidth per priority level.
- RSVP-TE control protocol packet counts.

When the optional **counters** keyword is specified, only the RSVP-TE control protocol packet counts are shown.

Example

The following command displays detailed RSVP-TE information for the interfaces associated with VLAN 1:

```
# show mpls rsvp-te interface vlan1 det
VLAN Name : vlan1
Local IP Address      : 11.121.96.20
MPLS Admin Status     : Enabled
MPLS Oper Status      : Enabled
RSVP-TE Admin Status  : Enabled
RSVP-TE Oper Status   : Enabled
RSVP-TE Up-Time       : 0d:1h:19m:46s
# Neighbors           : 1
Receive CIR           : 50000 Kbps
Transmit CIR          : 50000 Kbps
TE Metric             : Use IGP Cost/Metric
Hello Interval        : 3 seconds
Refresh Time          : 30 seconds
Hello Keep Multiplier : 3
Refresh Keep Multiplier : 3
Summary Refresh       : Disabled
Summary Refresh Time  : 3000 milliseconds
Bundle Message        : Disabled
Bundle Message Time   : 1000 milliseconds
```

Dir	Pool	CIR Available (per priority level)							
0	1	2	3	4	5	6	7		
Rx	50	50	50	50	50	50	50	50	50
Tx	50	50	50	50	50	50	50	50	40

```
(Rx)Receive Bandwidth (Tx)Transmit Bandwidth
Message                Sent      Received
-----
PATH                   165       165
PATH_TEAR              2         3
PATH_ERR               4        18
RESV                   160      145
RESV_TEAR              0         2
RESV_ERR               0         0
RESV_CONFIRM           4         1
SUMMARY_REFRESH        0         0
BUNDLE                 0         0
HELLO                  42        30
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

show mpls rsvp-te lsp

```
show mpls rsvp-te lsp {[destination | origin] ipaddress} {fast-reroute}
                        {detail} | summary}
```

Description

Displays complete or filtered information for all RSVP-TE LSPs.

Syntax Description

destination	Displays only those LSPs that terminate at the specified IP address.
origin	Displays only those LSPs that originate at the specified IP address.
<i>ipaddress</i>	Specifies an IP address for an LSP origin or destination.
fast-reroute	Limits the display to only those LSPs with fast reroute protection.
detail	Displays the LSP information in verbose format.
summary	Displays only the LSP summary statistics section of the normal display.

Default

N/A.

Usage Guidelines

If no options are specified, information for all RSVP-TE LSPs is displayed.

You can limit the display to ingress, transit, or egress LSPs with the following commands:

```
show mpls rsvp-te lsp [ingress {fast-reroute} | ingress_lsp_name
| ingressingress_lsp_name | ingress [destination | origin]ipaddress]
{[all-paths | detail] | summary | down-paths {detail}} show mpls rsvp-
te lsp [egress | transit] {fast-reroute} {[lsp_name] {[destination |
origin] ipaddress} {detail} | summary}
```

When label recording is enabled for an LSP, labels are displayed only for the ingress node (the egress label from the previous node always matches the ingress node label).

Example

The following command example displays information about all RSVP-TE LSPs:

```
# show mpls rsvp-te lsp
Ingress LSP Name Path Name      Destination      Next Hop I/F      UpTm  Flags
```



```

-----
frrlsp1tom2    any          11.100.100.50    m2vlan1          5m UEF---IV
lsptom2        pathtom2     11.100.100.50    m2vlan1          5m UES--OIV
tom5           any          11.100.100.55    m5vlan1          5m UEP---IV
Egress LSP Name Source IP      Destination      Prev Hop I/F      UpTm
-----
tom3           11.100.100.55    11.100.100.30    m5vlan1          5m
frrlsp1tom3    11.100.100.50    11.100.100.30    m2vlan1          5m
tom3           11.100.100.50    11.100.100.30    m2vlan1          5m
Transit LSP Name Source IP      Destination      Prev Hop I/F      Next Hop I/F      UpTm
-----
tom2           11.100.100.55    11.100.100.50    m5vlan1          m2vlan1          5m
frrlsp1tom5    11.100.100.50    11.100.100.55    m2vlan1          m5vlan1          5m
lsptom5        11.100.100.50    11.100.100.55    m2vlan1          m5vlan1          5m
Flags: (U) Up, (E) Enabled, (P) Primary LSP, (S) Secondary LSP,
(F) FRR Primary LSP (R) Redundant Paths, (B) Bandwidth Requested,
(O) ERO Specified, (I) IP Traffic Allowed, (V) VPN Traffic Allowed,
(v) VPN Assigned Traffic Allowed
Summary of RSVP-TE LSPs
Ingress LSPs (Enabled/Disabled)      : 3 (3/0)
Ingress LSPs with no configured path : 0
Ingress LSP Paths (Up/Down)          : 3 (3/0)
Detour LSP Paths (Up/Down)           : 0 (0/0)
Transit LSPs                         : 3
Egress LSPs                         : 3

```

The next command example displays only the summary information for all RSVP-TE LSPs:

```

# show mpls rsvp-te lsp summary
Summary of RSVP-TE LSPs
Ingress LSPs (Enabled/Disabled)      : 3 (3/0)
Ingress LSPs with no configured path : 0
Ingress LSP Paths (Up/Down)          : 3 (3/0)
Detour LSP Paths (Up/Down)           : 0 (0/0)
Transit LSPs                         : 3
Egress LSPs                         : 3

```

The following command example limits the display to RSVP-TE LSPs with fast reroute protection:

```

# show mpls rsvp-te lsp fast-reroute
Ingress LSP Name Path Name      Destination      Next Hop I/F      UpTm Flags
-----
frrlsp1tom2    any          11.100.100.50    m2vlan1          5m UEF---IV
Egress LSP Name Source IP      Destination      Prev Hop I/F      UpTm
-----
tom3           11.100.100.55    11.100.100.30    m5vlan1          5m
frrlsp1tom3    11.100.100.50    11.100.100.30    m2vlan1          5m
tom3           11.100.100.50    11.100.100.30    m2vlan1          5m
Transit LSP Name Source IP      Destination      Prev Hop I/F      Next Hop I/F      UpTm
-----
tom2           11.100.100.55    11.100.100.50    m5vlan1          m2vlan1          5m
frrlsp1tom5    11.100.100.50    11.100.100.55    m2vlan1          m5vlan1          5m
lsptom5        11.100.100.50    11.100.100.55    m2vlan1          m5vlan1          5m
Flags: (U) Up, (E) Enabled, (P) Primary LSP, (S) Secondary LSP,
(F) FRR Primary LSP (R) Redundant Paths, (B) Bandwidth Requested,
(O) ERO Specified, (I) IP Traffic Allowed, (V) VPN Traffic Allowed,
(v) VPN Assigned Traffic Allowed
Summary of RSVP-TE LSPs
Ingress FRR LSPs (Enabled/Disabled)   : 1 (1/0)
Ingress FRR LSPs with no configured path : 0
Ingress FRR LSP Paths (Up/Down)       : 1 (1/0)
Detour LSP Paths (Up/Down)            : 0 (0/0)

```

```
Transit FRR LSPs           : 3
Egress FRR LSPs           : 3
```

The following command example displays detailed information for all RSVP-TE LSPs:

```
# show mpls rs lsp det

Ingress LSP Name: lsp_to_baha
  Destination : 11.100.100.4      Admin Status : Enabled
  IP Traffic  : Allow             #VPLS Cfgd   : 0
  VPN Traffic : Allow             #VPLS In-Use : 0

  Path Name: path1
    Oper Status : Enabled          UpTime           : 0d:0h:4m:49s
    Profile Name : default
    Peak Rate    : 0 Kbps          Max Burst Size   : 0 Kb
    Committed Rate : 0 Kbps        Setup/Hold Priority : 7/0
    Record Route : Enabled
    MTU          : Use Local I/F
    Tunnel ID    : 1              Ext Tunnel ID     : 11.100.100.1
    LSP ID       : 0              State Changes      : 1
    LSP Type     : Primary         Bandwidth Cfgd    : False
    Activity     : Active
    Failures     : 0              Retries-since last failure : 0
    Retries-Total : 0
    Configured ERO : Order  IP Address/Mask  Type  Inc/Exc
                   100  9.50.1.2/32          strict include
    Advertised Label: n/a          Received Label : 0x00434
    Rx Packets      : n/a          Tx Packets      : --
    Rx Bytes       : n/a          Tx Bytes       : --
    Next Hop I/F   : 9.50.1.1 - j1-j2vlan1
    Next Hop Addr  : 9.50.1.2
    Record Route   : Indx  IP Address      Label
                   :    1  9.50.1.2        0x00434
                   :    2  9.54.1.4        0x00434

  Path Name: path2
    Oper Status : Enabled          UpTime           : 0d:0h:4m:49s
    Profile Name : default
    Peak Rate    : 0 Kbps          Max Burst Size   : 0 Kb
    Committed Rate : 0 Kbps        Setup/Hold Priority : 7/0
    Record Route : Enabled
    MTU          : Use Local I/F
    Tunnel ID    : 2              Ext Tunnel ID     : 11.100.100.1
    LSP ID       : 0              State Changes      : 1
    LSP Type     : Secondary       Bandwidth Cfgd    : False
    Activity     : Standby
    Failures     : 0              Retries-since last failure : 0
    Retries-Total : 0
    Configured ERO : Order  IP Address/Mask  Type  Inc/Exc
                   100  9.50.1.2/32          n/a      exclude
    Advertised Label: n/a          Received Label : 0x00435
    Rx Packets      : n/a          Tx Packets      : --
    Rx Bytes       : n/a          Tx Bytes       : --
    Next Hop I/F   : 9.52.1.1 - j1-j4vlan1
    Next Hop Addr  : 9.52.1.4
    Record Route   : Indx  IP Address      Label
                   :    1  9.52.1.4        0x00435
```

History

This command was first available in ExtremeXOS 11.6.

This command and its output were modified, and the summary option was added in ExtremeXOS 12.0.

The fast-reroute feature was first available in ExtremeXOS 12.1.

This command and its output were modified in ExtremeXOS 12.2.2.

The output was modified to include the include/exclude information in ExtremeXOS 15.7.1.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

show mpls rsvp-te lsp [egress | transit]

```
show mpls rsvp-te lsp [egress | transit] {fast-reroute} {{lsp_name}
  { [destination | origin] ipaddress } {detail} | summary}
```

Description

Displays complete or filtered information for one or all egress or transit RSVP-TE LSPs.

Syntax Description

egress	Limits the display to only the LSPs that terminate at this switch.
transit	Limits the display to only the LSPs that transit this switch.
fast-reroute	Limits the display to only those LSPs with fast reroute protection.
<i>lsp_name</i>	When either the transit or the egress option is specified, this variable specifies a name for a single LSP for which information is displayed.
destination	Displays only those LSPs that terminate at the specified IP address.
origin	Displays only those LSPs that originate at the specified IP address.
<i>ipaddress</i>	Specifies an IP address for an LSP origin or destination.
detail	Displays the LSP information in verbose format.
summary	Displays only the LSP summary statistics section of the normal display.

Default

N/A.

Usage Guidelines

You can limit the display to ingress LSPs with the following command:

```
show mpls rsvp-te lsp [ingress {fast-reroute} | ingress_lsp_name
| ingress ingress_lsp_name | ingress [destination | origin] ipaddress]
{[all-paths | detail] | summary | down-paths {detail}}
```

You can display information for all LSPs with the following command:

```
show mpls rsvp-te lsp {[destination | origin] ipaddress} {fast-reroute}
{detail} | summary
```

When label recording is enabled for an LSP, labels are displayed only for the ingress node (the egress label from the previous node always matches the ingress node label).

Example

The following command example displays RSVP-TE LSPs that terminate at this switch and at IP address 11.100.100.30:

```
# show mpls rsvp-te lsp egress destination 11.100.100.30
Egress LSP Name      Source IP      Destination    Prev Hop I/F      UpTm
-----
tom3                 11.100.100.55  11.100.100.30  m5vlan1           5m
frrlsp1tom3         11.100.100.50  11.100.100.30  m2vlan1           5m
tom3                 11.100.100.50  11.100.100.30  m2vlan1           5m
Summary of Egress RSVP-TE LSPs to destination 11.100.100.30
Egress LSPs                      : 3
Egress Protected LSPs           : 0
```

History

This command was first available in ExtremeXOS 11.6.

This command and its output were modified, and the summary option was added in ExtremeXOS 12.0.

The fast-reroute feature was first available in ExtremeXOS 12.1.

This command and its output were modified in ExtremeXOS 12.2.2.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

show mpls rsvp-te lsp ingress

```
show mpls rsvp-te lsp [ingress {fast-reroute} | ingress_lsp_name |
    ingress ingress_lsp_name | ingress [destination | origin] ipaddress]
    {[all-paths | detail] | summary | down-paths {detail}}
```

Description

Displays information for the specified ingress RSVP-TE LSP.

Syntax Description

fast-reroute	Limits the display to only those LSPs with fast reroute protection.
<i>ingress_lsp_name</i>	Identifies the ingress LSP for which you want to display information.
destination	Displays only those LSPs that terminate at the specified IP address.
origin	Displays only those LSPs that originate at the specified IP address.
<i>ipaddress</i>	Specifies an IP address for an LSP origin or destination.
all-paths	Specifies that the display include all redundant paths.
detail	Displays the LSP information in verbose format.
summary	Displays only the LSP summary statistics section of the normal display.
down-paths	Specifies that the display include only those paths that are operationally down.

Default

N/A.

Usage Guidelines

You can limit the display to egress or transit LSPs with the following command:

```
show mpls rsvp-te lsp [egress | transit] {fast-reroute} {{lsp_name}
[destination | origin]ipaddress} {detail} | summary}
```

You can display information for all LSPs with the following command:

```
show mpls rsvp-te lsp {[destination | origin] ipaddress} {fast-reroute}
{detail} | summary}
```

When label recording is enabled for an LSP, labels are displayed only for the ingress node (the egress label from the previous node always matches the ingress node label).

Example

Use the following command to display information about a specific LSP:

```
# show mpls rsvp-te lsp jefflsp1
Ingress LSP Name Path Name      Destination      Next Hop I/F      UpTm  Flags
-----
jefflsp1      jeffpath1      11.100.100.204   n/a              0    --PR--IV
Flags: (U) Up, (E) Enabled, (P) Primary LSP, (S) Secondary LSP,
(F) FRR Primary LSP (R) Redundant Paths, (B) Bandwidth Requested,
(O) ERO Specified, (I) IP Traffic Allowed, (V) VPN Traffic Allowed,
(v) VPN Assigned Traffic Allowed
Summary of Ingress RSVP-TE LSPs named jefflsp1
Ingress LSPs (Enabled/Disabled)      : 1 (0/1)
Ingress LSPs with no configured path : 0
Ingress LSP Paths (Up/Down)          : 2 (0/2)
```

Use the following command to display detailed information about a specific ingress LSP:

```
# show mpls rsvp-te lsp ingress "lsp598" detail
Ingress LSP Name: lsp598
```

```

Destination : 11.100.100.8      Admin Status : Enabled
IP Traffic   : Allow            #VPLS Cfgd   : 0
VPN Traffic  : Allow            #VPLS In-Use : 0
Path Name: path598
Profile Name : prof598
Tunnel ID    : 1                Ext Tunnel ID : 11.100.100.20
LSP ID       : 0                State Changes : 5
Oper Status  : Enabled          Bandwidth Cfgd : False
LSP Type     : Primary
Activity     : Active
Failures     : 2                Retries-since last failure : 0
Retries-Total : 12
Rcv Label    : 0x0052e          UpTime       : 0d:0h:3m:44s
Next Hop     : 11.121.96.5
Tx I/F       : 11.121.96.20 - vlan1
Record Route : Indx IP Address
: 1 11.121.96.5
: 2 11.95.96.9
: 3 11.98.96.8

```

Use the following command to display detailed information about all paths for an LSP:

```

# show mpls rsvp-te lsp jefflsp1 all-paths
Ingress LSP Name Path Name      Destination      Transmit I/F      UpTm  Flags
-----
jefflsp1          jeffpath0        11.100.100.204   n/a              0  --SR--IV
jefflsp1          jeffpath1        11.100.100.204   n/a              0  --PR--IV
Flags: (U) Up, (E) Enabled, (P) Primary LSP, (S) Secondary LSP,
(F) FRR Primary LSP (R) Redundant Paths, (B) Bandwidth Requested,
(O) ERO Specified, (I) IP Traffic Allowed, (V) VPN Traffic Allowed,
(v) VPN Assigned Traffic Allowed
Summary of Ingress RSVP-TE LSPs named jefflsp1
Ingress LSPs (Enabled/Disabled)      : 1 (0/1)
Ingress LSPs with no configured path : 0
Ingress LSP Paths (Up/Down)          : 2 (0/2)

```

Use the following command to display information about all ingress down paths:

```

# show mpls rsvp-te lsp ingress down-paths
Ingress LSP Name Path Name      Destination      Transmit I/F      UpTm  Flags
-----
jefflsp1          jeffpath0        11.100.100.204   n/a              0  --SR--IV
jefflsp1          jeffpath1        11.100.100.204   n/a              0  --PR--IV
jefflsp2          jeffpath2        11.100.100.203   n/a              0  -EPR-OIV
Flags: (U) Up, (E) Enabled, (P) Primary LSP, (S) Secondary LSP,
(F) FRR Primary LSP (R) Redundant Paths, (B) Bandwidth Requested,
(O) ERO Specified, (I) IP Traffic Allowed, (V) VPN Traffic Allowed,
(v) VPN Assigned Traffic Allowed
Summary of Ingress RSVP-TE LSPs
Ingress LSP Paths that are Down      : 3

```

History

This command was first available in ExtremeXOS 11.6.

This command and its output were modified, and the summary option was added in ExtremeXOS 12.0.

The command output was modified in ExtremeXOS 12.2.2.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

show mpls rsvp-te neighbor

```
show mpls rsvp-te neighbor [{vlan} vlan_name | ipaddress] {detail}
```

Description

Displays all recognized RSVP-TE neighbors.

Syntax Description

<i>vlan_name</i>	Displays only the neighbors for the specified VLAN .
<i>ipaddress</i>	Displays only the neighbor with the specified ipaddress.

Default

N/A.

Usage Guidelines

This command displays all recognized RSVP-TE neighbors. The IP address of each neighbor is displayed along with the VLAN name for the [MPLS](#) interface. For each neighbor, the following information is displayed:

- Number of RSVP-TE LSPs.
- Number of hello periods that have elapsed without receiving a valid hello.
- Remaining time before next hello is sent.
- Remaining time before next bundle message is sent.
- Neighbor up time.
- Neighbor supports RSVP hello.
- RSVP hello state.
- Neighbor supports refresh reduction.

If *vlan_name* is specified, only neighbors for the matching VLAN are shown. If *ipaddress* is specified, only the neighbor with that IP address is shown. If the **detail** keyword is specified, the information is shown in a verbose manner.

Example

The following command displays all recognized RSVP-TE neighbors:

```
# show mpls rsvp-te neighbor
NeighborIP      VLAN Name      #LSPs #Miss NxtHello NxtBundl Flag UpTm
-----
```

```

11.121.96.5      vlan1          2      0      2870      0 UTH-  15m
11.122.96.8      vlan2          2      0      1770      0 UTH-  15m
Flags: (U) Hello Session Up, (T) Two Way Hello, (H) Neighbor Supports Hello,
(R) Neighbor Supports Refresh Reduction

```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

show mpls rsvp-te path

```
show mpls rsvp-te path {path_name} {detail}
```

Description

Displays the configuration and usage information for MPLS RSVP-TE routed paths.

Syntax Description

<i>path_name</i>	Displays configuration and usage information for the specified MPLS RSVP-TE path.
detail	Displays the path information in verbose format.

Default

N/A.

Usage Guidelines

This command displays the configuration and usage information for MPLS RSVP-TE paths. Information is listed in tabular format and includes:

- Path name.
- Number of configured ERO objects.
- Number of LSPs configured to use this path.
- List of EROs and their type.
- Whether the hop is included in the path calculation.

Specifying the optional **detail** keyword displays the path information in verbose format. If **detail** is specified, all LSPs that are configured to use the path are also displayed.

Example

The following example displays configuration and status information for the MPLS RSVP-TE paths:

```
# show mpls rs path
Path Name      #LSP #ERO Ord#  ERO IP Netmask      Type  Inc/Exc
-----
path1          1    1   100 9.50.1.2/32        strict include
path2          1    1   100 9.50.1.2/32        n/a    exclude
```

The following example displays configuration and status information for the MPLS RSVP-TE paths in the verbose format:

```
# show mpls rs path det

Path Name : path1
Hop List Index      : 1
Path Option Index   : 1
#ERO/Hops           : 1
#LSP References     : 1
ERO                 : Order#  IPAddress/Mask      Type  Inc/Exc
                   : -----
                   :      100 9.50.1.2/32        strict include

Path Name : path2
Hop List Index      : 2
Path Option Index   : 1
#ERO/Hops           : 1
#LSP References     : 1
ERO                 : Order#  IPAddress/Mask      Type  Inc/Exc
                   : -----
                   :      100 9.50.1.2/32        n/a    exclude
```

History

This command was first available in ExtremeXOS 11.6.

The output was modified to include the include/exclude information in ExtremeXOS 15.7.1.

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

show mpls rsvp-te profile

```
show mpls rsvp-te profile {profile_name} {detail}
```

Description

Displays the configuration for the specified profile.

Syntax Description

<i>profile_name</i>	Displays configuration and usage information for the specified profile, which can be either a standard or a fast-reroute profile.
detail	Displays the profile information in verbose format, and displays all LSPs that are configured to use the specified profile.

Default

N/A.

Usage Guidelines

If the *profile_name* argument is omitted, the profile parameter values for all profiles are displayed.

Example

The following command displays configuration information for all defined profiles:

```
# show mpls rsvp-te profile
Profile Name      Peak      Committed  Max Burst  SPri  HPri  RRO  MTU  #LSP
-----
default          0          0          0         7     0    Off i/f    0
prof598          0          0          0         7     0    On  1500    1
FRR Profile Name Mode    Bandwidth  SPri  HPri  HopLmt  P-BW  P-Node  #LSP
-----
Prfl_frr         Detour      0         7     0         3    Ena   Ena      2
```

The following command displays configuration information for a specific fast-reroute profile:

```
# show mpls rsvp-te profile prfl_frr detail
Profile Name : prfl_frr
Profile type  : Fast Reroute / Standard
Peak Rate    : 0 Kbps
Committed Rate : 0 Kbps
Max Burst Size : 0 Kb
Setup Priority : 7
Hold Priority  : 4
Hop Limit     : 1
Protected BW  : Disabled
Protect Node   : Enabled
#LSP References : 2
LSP / Path    : pc10_lsp / p1
```

History

This command was first available in ExtremeXOS 11.6.

The fast-reroute feature was first available in ExtremeXOS 12.1.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

show mpls rsvp-te profile fast-reroute

```
show mpls rsvp-te profile fast-reroute {detail}
```

Description

Displays the configuration for all fast-reroute profiles.

Syntax Description

detail	Displays the profile information in verbose format.
---------------	---

Default

N/A.

Usage Guidelines

None.

Example

The following command displays summary configuration information for all fast-reroute profiles:

```
# show mpls rsvp-te profile fast-reroute
FRR Profile Name Mode   Bandwidth SPri HPri HopLmt P-BW P-Node #LSP
-----
default_frr      Detour      10      7    0      3 Dis  Ena     4
prfl_frr         Detour       0      7    0      3 Dis  Ena     2
```

The following command displays detailed configuration information for all fast-reroute profiles:

```
# show mpls rsvp-te profile fast-reroute detail
Profile Name : prfl_frr
Peak Rate    : 0 Kbps
Committed Rate : 0 Kbps
Max Burst Size : 0 Kb
Setup Priority : 7
Hold Priority : 4
Hop Limit    : 1
Protected BW : Off
Protect Node  : On
#LSP References : 2
LSP / Path   : pc10_lsp / p1
Profile Name : default_frr
Peak Rate    : 0 Kbps
Committed Rate : 0 Kbps
Max Burst Size : 0 Kb
```

```
Setup Priority      : 4
Hold Priority       : 0
Hop Limit          : 2
Protected BW       : On
Protect Node       : On
#LSP References    : 4
LSP / Path         : pc_frr_lsp / p2
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

show mpls rsvp-te

```
show mpls rsvp-te
```

Description

Displays displays summary configuration and status information for RSVP-TE.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command displays summary configuration and status information for RSVP-TE. The parameters displayed include:

- Global status of RSVP-TE.
- Configured standard LSP timer values.
- Configured rapid-retry LSP timer values.
- RSVP-TE VLANs.

Example

The following command shows the summary configuration and status information for RSVP-TE:

```
# show mpls rsvp-te
RSVP-TE Admin Status : Enabled
RSVP-TE Oper Status  : Enabled
```

```
LSP Standard-Retry Timers
Delay-Interval   : 30 seconds
Decay-Rate       : 50 %
Retry-Limit      : unlimited
LSP Rapid-Retry Timers
Delay-Interval   : 500 milliseconds
Decay-Rate       : 50 %
Retry-Limit      : 10
RSVP-TE VLANs   : vlan1
                  : vlan2
                  : vlan3
```

History

This command was first available in ExtremeXOS 11.6

Platform Availability

This command is available only on the platforms that support *MPLS* as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

show mpls static lsp

```
show mpls static lsp {summary | {lsp_name} {detail}}
```

Description

Displays the configuration of one or all static LSPs.

Syntax Description

summary	Displays only static LSP summary information.
<i>lsp_name</i>	Identifies the LSP to be displayed.
detail	Displays additional information about the static LSPs, including packet and byte counts.

Default

N/A.

Usage Guidelines

If no command options are specified, all defined static LSPs are displayed in tabular format. The information displayed includes the configured ingress label, egress label, next-hop router IP address, and the *MPLS* interface status for the egress path. The summarized list of static LSPs is displayed in alphabetical order based on the LSP name.

Example

The following command displays detailed information about an ingress static LSP:

```
show mpls static lsp statlsp1 detail
Static LSP Name : statlsp1          LSP Type      : Ingress-Lsp
Destination   : 11.100.100.55      Next Hop Addr : 12.220.0.30
Ingress I/F   : **None**           Egress I/F    : m3vlan1
Admin Status  : Disabled            Oper Status   : Disabled
Ing-Label     : **None**           Eg-Label      : 0x7FF00
IP Traffic    : Allow               VPN Traffic   : Allow
Rx Packets    : n/a                 Tx Packets    : 0
Rx Bytes      : n/a                 Tx Bytes      : 0
```

The following command displays detailed information about a transit static LSP:

```
show mpls static lsp statlsp1 detail
Static LSP Name : statlsp1          LSP Type      : Transit-Lsp
Destination   : 11.100.100.55      Next Hop Addr : 12.224.0.55
Ingress I/F   : m2vlan1            Egress I/F    : m5vlan1
Admin Status  : Disabled            Oper Status   : Disabled
Ing-Label     : 0x7FF00             Eg-Label      : 0x80300
IP Traffic    : Not Applicable      VPN Traffic   : Not Applicable
Rx Packets    : 0                   Tx Packets    : 0
Rx Bytes      : 0                   Tx Bytes      : 0
```

The following command displays detailed information about an egress static LSP:

```
show mpls static lsp statlsp1 detail
Static LSP Name : statlsp1          LSP Type      : Egress-Lsp
Destination   : 11.100.100.55      Next Hop Addr : **None**
Ingress I/F   : m3vlan1            Egress I/F    : **None**
Admin Status  : Disabled            Oper Status   : Disabled
Ing-Label     : 0x80300             Eg-Label      : **None**
IP Traffic    : Not Applicable      VPN Traffic   : Not Applicable
Rx Packets    : 0                   Tx Packets    : n/a
Rx Bytes      : 0                   Tx Bytes      : n/a
```

History

This command was first available in ExtremeXOS 12.1.

The output for this command was modified in ExtremeXOS 12.2.2.

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

show mpls statistics l2vpn

```
show mpls statistics l2vpn {vpls_name | vpws_name} {detail}
```

Description

Displays *MPLS* statistics for one or all Layer 2 VPNs.

Syntax Description

<i>vpls_name</i>	Specifies VPLS for which to display statistics.
<i>vpws_name</i>	Specifies VPWS for which to display statistics.
detail	Displays additional information about the Layer 2 VPNs.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays statistics for all Layer 2 VPNs:

```
Switch.1 # show mpls statistics l2vpn
VPN ID      Peer IP              RxPackets      RxBytes      TxPackets      TxBytes
-----
99          11.100.100.219              32866         4967734      14005          2394407
11.100.100.218          398             8235          577          10583
2009        11.100.100.219              0              0              0              0
2008        11.100.100.219              0              0              15             688
```

The following command displays detailed statistics for all Layer 2 VPNs:

```
Switch.2 # (debug) 12.2 # show mpls statistics l2vpn detail
VPNID (L2VPN Name)
Peer IP      State RxLabel TxLabel LSPTxLabel NextHopI/F
RxPackets    RxBytes TxPackets TxBytes
-----
99 (jwcvpls)
11.100.100.219 Up    x80402 x80402 x00010 tordoze
32866         4967734          14005          2394407
11.100.100.218 Up    x80407 x80405 x00011 tornext
398           8235           577           10583
2009 (pws-1)
11.100.100.219 Up    x80403 x80403 x00010 tordoze
0             0              0              0
2008 (pws-2)
11.100.100.219 Up    x80404 x80404 x00010 tordoze
0             0              15             688
```

The following show output is for ExtremeXOS 15.4:

```
L2VPN Name
Peer IP/LSP  Flags      RxPackets  RxBytes    TxPackets  TxBytes
-----
vpls1
```

```

1.1.1.1      U          0          0          0          0
  lsp1       +          0          0          0          0
  lsp2       +          0          0          0          0

```

Flags: (U) Up, (D) Down, (R) Ready, (S) Signaling, (+) In-use,
(-) Not In-use

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

show mrp ports

```
show mrp ports {port_list}
```

Description

Shows the MRP timers configured on the given list of ports on the switch.

Syntax Description

mrp	Multiple Registration Protocol.
<i>port_list</i>	Ports on which MRP timers are configured or unconfigured.

Default

N/A.

Usage Guidelines

Use this command to view MRP timers configured on the given list of ports on the switch.

Example

```

# show mrp ports 1, 4, 5
-----
Ports      Join Time (ms)      Leave Time (ms)      Leave All Time
(ms)      Periodoc (ms)      Extended
-----
Refresh (ms)
-----
1          200          600

```


10000	1000	10000
4	300	800
10000	1000	10000
5	200	600
10000	1000	10000

History

This command was first available in ExtremeXOS 15.3.

Output for periodic and extended refresh timers added in 15.3.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show msdp memory

show msdp memory {**detail** | *memoryType*}

Description

This command displays current memory utilization of the MSDP process, including all virtual router instances of the MSDP process.

Syntax Description

detail	Displays detailed statistics for all memory types.
<i>memoryType</i>	Displays statistics for a particular memory type.

Default

N/A.

Usage Guidelines

Use this command to view and diagnose the memory utilization of the MSDP process.

Example

The following displays current memory utilization of the MSDP process, including all virtual router instances of the MSDP process:

```
show msdp memory
```

The following is sample output from this command:

```
MSDP Memory Information
-----
```

Bytes Allocated: 79792 AllocFailed: 0 OversizeAlloc: 0
Current Memory Utilization Level: GREEN
Memory Utilization Statistics

Size	16	32	48	64	80	96	128	256	1024	4096	8192	12288
-----	Used	Blocks	0	0	256	263	3	0	2	0	1	
0	0	4										
peer	0	0	0	0	0	0	0	0	0	0	0	4
mesh-group	0	0	0	0	3	0	0	0	0	0	0	
0	0											
sa-node	0	0	0	255	0	0	0	0	0	0	0	0
sa-entry	0	0	255	0	0	0	0	0	0	0	0	0
0												
vr-node	0	0	0	0	0	0	0	0	0	1	0	0
rt-cache	0	0	0	5	0	0	0	0	0	0	0	0
0												
rp-node	0	0	1	0	0	0	0	0	0	0	0	0
client	0	0	0	0	0	0	2	0	0	0	0	0
misc	0	0	0	0	3	0	0	0	0	0	0	0

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

show msdp mesh-group

```
show msdp [mesh-group {detail} | {mesh-group} mesh-group-name] {vr  
  vrname}
```

Description

This command displays configuration information about MSDP mesh-groups.

Syntax Description

detail	Displays detailed information about MSDP mesh-groups.
<i>mesh-group-name</i>	Specifies the name of the MSDP mesh-group. The character string can be a maximum of 31 characters.
<i>vrname</i>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

N/A.

Usage Guidelines

Use this command to display configuration information about MSDP mesh-groups, as follows:

- For summary information, enter the `show mdp mesh-group` command.
- For detailed information, enter the `show mdp mesh-group detail` command.
- For detailed information about a specific mesh-group, enter the `show mdp mesh-group name` command.

Example

The following command displays the peer count for a mesh-group:

```
show mdp mesh-group
```

The following is sample output from this command:

MeshGroupName	PeerCount
external	0
internal	0
mdp_mesh	4

The following command displays detailed information about a mesh-group called "mdp_mesh":

```
show mdp mesh-group "mdp_mesh"
```

The following is sample output from this command:

Mesh Group Name	:	mdp_mesh	Num of Peers	:	4
Peers	:	54.172.168.97	55.0.0.83	:	124.56.78.90
221.160.90.228	:			:	

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

show mdp peer

```
show mdp [peer {detail} | {peer} remoteaddr] {vr vr_name}
```

Description

This command displays configuration and run-time parameters about *MSDP* peers.

Syntax Description

detail	Displays detailed information about MSDP peers.
<i>remoteaddr</i>	Specifies the IP address of the MSDP peer.
<i>vr_name</i>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

N/A.

Usage Guidelines

Use this command to verify the configuration and run-time parameters for MSDP peers, as follows:

- For summary information, enter the `show msdp peer` command.
- For detailed information for all peers, enter the `show msdp peer detail` command.
- For detailed information for a specific peer, enter the `show msdp peer remoteaddr` command.

Example

The following command displays configuration and run-time parameters for MSDP peers:

```
show msdp peer
```

The following is sample output from this command:

```

Peer Address      AS      State      Up/Down    Resets    SA_Cnt    Name
-----
-d 54.172.168.97   14490   DISABLED   00:31:36   0         0         test
*e 55.0.0.83       100     ESTABLISHED 00:21:04   1         0         to-Hawaii
-d 124.56.78.90    2345    DISABLED   00:31:36   0         0
-d 221.160.90.228  23456   DISABLED   00:31:36   0         0
Flags: (*) default peer, (d) disabled, (e) enabled

```

The following command displays output from an MSDP peer with the IP address 16.0.0.2:

```

* Switch.8 # show msdp peer 16.0.0.2
MSDP Peer           : 16.0.0.2
Enabled             : No
Keepalive Interval  : 60
Source Address      : not known
Default Peer        : No
Process In Request  : Yes
Maximum SA Limit    : not configured
Input SA Filter     : not configured
State               : DISABLED
Local Port          : 0
In Total Msgs       : 0
In SA Msgs          : 0
In SA Req Msgs      : 0
In SA Resp Msgs     : 0
Time since Last Msg : 00:00:02
Connection Attempts : 0
RPF Fails           : 0
AS Number           : 100.100
Holdtimer Interval  : 75
TTL Threshold       : 0
Default Peer Filter : not configured
In Request filter   : not configured
Mesh Group          : not configured
Output SA Filter    : not configured
Uptime/Downtime     : 00:00:02
Remote Port         : 0
Out Total Msgs      : 0
Out SA Msgs         : 0
Out SA Req Msgs     : 0
Out SA Resp Msgs    : 0
Hold Tmr Exp in     : 00:00:00
Entered Established  : 0
Output Queue Size   : 0

```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

show msdp sa-cache

```
show msdp [sa-cache | rejected-sa-cache] {group-address grp-addr {source-address src-addr {as-number as-num} {originator-rp originator-rp-addr} {local} {peer remoteaddr} {vr vrname}
```

Description

This command displays the SA cache database. The following quadruplet per SA cache entry displays: {Group, Source, originating RP, and peer}. In addition, information about the following displays: the cache uptime, aging, whether sources are local or remote, etc.

Syntax Description

<i>grp-addr</i>	Displays the SA cache within the specified group address range.
<i>src-addr</i>	Displays the SA cache within the specified source address range.
<i>as-num</i>	Displays all SA cache that originated from the specified Autonomous System (AS) number.
<i>originator-rp-addr</i>	Displays all SA cache entries that were originated by the specified rendezvous point.
local	Displays locally originated SA cache entries only.
<i>remoteaddr</i>	Displays the SA cache entries received from the <i>MSDP</i> peer with the specified IP address.
<i>vrname</i>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.

Default

N/A.

Usage Guidelines

Use this command to view and troubleshoot the SA cache database. There are various filtering criteria you can use to display just a subset of the SA cache database. The following are some of the criteria, which you can use together or separately, to display information about the SA cache:

- Filtering on the group address range.
- Filtering on the source address range.
- Filtering on the originator rendezvous point address.
- Filtering of the advertising MSDP peer.
- Locally originated SA cache.
- Rejected SA cache.

Example

The following command displays the SA cache database:

```
# show msdp sa-cache
Group Address      Source Address    Originator        Peer Address      Age/Ageout In
-----
235.100.200.1      10.20.30.1       60.0.0.5         10.0.0.1          00:44:24/05:10
235.100.200.2      10.20.30.2       60.0.0.5         192.0.0.16        00:44:24/05:16
235.100.200.3      10.20.30.3       60.0.0.5         10.0.0.1          00:44:24/05:10
235.100.200.4      10.20.30.4       60.0.0.5         10.0.0.1          00:44:24/05:10
235.100.200.5      10.20.30.5       60.0.0.5         55.0.0.5          00:44:24/05:01
235.100.200.6      10.20.30.6       60.0.0.5         178.54.67.23      00:44:24/05:17
235.100.200.7      10.20.30.7       60.0.0.5         112.234.213.12    00:44:24/05:43
235.100.200.8      10.20.30.8       60.0.0.5         10.0.0.1          00:44:24/05:10
235.100.200.9      10.20.30.9       60.0.0.5         10.0.0.1          00:44:24/05:10
235.100.200.10     10.20.30.10      60.0.0.5         0.0.0.0           00:44:24/00:00
235.100.200.11     10.20.30.11      60.0.0.5         0.0.0.0           00:44:24/00:00
235.100.200.12     10.20.30.12      60.0.0.5         0.0.0.0           00:44:24/00:00
235.100.200.13     10.20.30.13      60.0.0.5         0.0.0.0           00:44:24/00:00
235.100.200.14     10.20.30.14      60.0.0.5         0.0.0.0           00:44:24/00:00
235.100.200.15     10.20.30.15      60.0.0.5         0.0.0.0           00:44:24/00:00
235.100.200.16     10.20.30.16      60.0.0.5         0.0.0.0           00:44:24/00:00
235.100.200.17     10.20.30.17      60.0.0.5         0.0.0.0           00:44:24/00:00
235.100.200.18     10.20.30.18      60.0.0.5         0.0.0.0           00:44:24/00:00
235.100.200.19     10.20.30.19      60.0.0.5         0.0.0.0           00:44:25/00:00
Number of accepted SAs      : 255
Number of rejected SAs      : 0
Flags: (a) Accepted, (f) Filtered by policy, (r) RPF check failed
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

show msdp

show msdp {**vr** *vrname*}

Description

This command displays global configuration and run-time parameters for *MSDP*.

Syntax Description

<i>vrname</i>	Specifies the name of the virtual router to which this command applies. If a name is not specified, it is extracted from the current CLI context.
---------------	---

Default

N/A.

Usage Guidelines

Use this command to verify the global configuration parameters of MSDP.

Example

The following command displays global configuration and run-time parameters for MSDP:

Switch.2 # show msdp			
MSDP Enabled	: No	VR-Name	: VR-Default
Originator RP Addr	: not configured	SA Cache ageout time	: 360
Store SA Cache	: Yes	SA Cache Server	: not configured
Export Local SAs	: Yes	Export SA filter	: not configured
Max Rejected Cache	: not configured	Encapsulate data	: Yes
Num of Rejected SAs	: 0	Total Num of SAs	: 0
Num of Local SAs	: 0	AS Disp Format	: Asdot

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

show msrp

show msrp

Description

Displays the MSRP configuration on the switch.

Syntax Description

msrp	Multiple Stream Registration Protocol.
-------------	--

Default

N/A.

Usage Guidelines

Use this command to display MSRP configuration on the switch.

Example

```
# show msrp
MSRP Status                : Enabled
MSRP Max Latency Frame Size : 1522
MSRP Max Fan-in Ports      : No limit
MSRP First Value Change Recovery Time : 10000 (ms)
MSRP Ignore Latency Changes : On
MSRP Talker VLAN Pruning   : On
MSRP Enabled Ports         : *17ab      *19a      !5
Total MSRP streams          : 4
Total MSRP reservations     : 2
Flags:      (*) Active,      (!) Administratively disabled,
            (a) SR Class A allowed, (b) SR Class B allowed.
```

History

This command was first available in ExtremeXOS 15.3.

The MSRP First Value Change Recovery Time, MSRP Ignore Latency Change, and MSRP Talker VLAN Pruning example outputs were added in 15.3.2.

Platform Availability

This command is available on all platforms that support the AVB feature and that have an AVB feature pack license installed. To see which platforms support AVB and for information about obtaining a license, see the [Switch Engine 32.3 Feature License Requirements](#).

show msrp listeners

```
show msrp listeners {egress | ingress | ingress-and-egress} {port
    port_num} {source-mac-addr source_mac_addr | stream-id stream_id}
```


Description

Shows MSRP listener information.

Syntax Description

msrp	Multiple Stream Registration Protocol.
listeners	Listener attributes.
egress	Display egress listeners only.
ingress	Display ingress listeners only (default).
ingress-and-egress	Display all listeners.
<i>port_num</i>	Filter based on ingress port number of the stream.
source-mac-addr	Filter based on source MAC address of a data stream.
stream-id	Filter based on stream ID of a data stream.

Default

N/A.

Usage Guidelines

Use this command to show MSRP listener information. The output can be filtered based on the stream id, source MAC, or port number on which the listener is registered.

Example

```
# show msrp listeners
      Stream Id          Port  Dec      Dir      State      Stream Age
                                     App  Reg  (days, hr:mm:ss)
-----
00:50:c2:4e:d3:2d:00:00    19  Ready  Ingress  VO    IN        0, 00:58:12
00:50:c2:4e:d3:2d:00:01    19  Ready  Ingress  VO    IN        0, 00:58:12
00:50:c2:4e:d3:2d:00:02    19  Ready  Ingress  VO    IN        0, 00:58:12
-----
App      : Applicant State,
Dir      : Direction of MSRP attribute,
Dec      : MSRP Declaration Types,
Reg      : Registrar State

MSRP Declaration Types:
AskFail  : Listener Asking Failed,
Ready    : Listener Ready
RdyFail  : Listener Ready Failed,

Applicant States:
AA       : Anxious active,
AO       : Anxious observer,
LA       : Leaving active,
QA       : Quiet active,
QP       : Quiet passive,
VO       : Very anxious observer,
AN       : Anxious new,
AP       : Anxious passive,
LO       : Leaving observer,
QO       : Quiet observer,
VN       : Very anxious new,
VP       : Very anxious passive

Registrar States:
```

IN : In - Registered, LV : Leaving - Timing out,
MT : Empty - Not Registered

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on all platforms that support the AVB feature and that have an AVB feature pack license installed. To see which platforms support AVB and for information about obtaining a license, see the [Switch Engine 32.3 Feature License Requirements](#).

show msrp ports

show msrp ports {port_list}{detail}

Description

Displays the MSRP configured port information.

Syntax Description

msrp	Multiple Stream Registration Protocol.
ports	Ports.
<i>port_list</i>	Port list separated by a comma or "-".
detail	Port information with more detail.

Default

N/A.

Usage Guidelines

Displays the MSRP configured port information. Specifying **detail** displays port information with more detail.

Example

```
# show msrp ports
```

Port	Enabled	Oper	Port	Dplx	Jumbo	Jumbo	Cls	Bndry	State	Sr-Pvid
			Speed			Size			App/Reg	
----	-----	-----	-----	----	-----	-----	---	-----	-----	-----
5	Y	Up/dbg			N	9216	A	N	QA/IN	2
							B	N	QA/IN	2

```

*21 Y Up 1000 M Full N 9216 A N QA/IN 2
                                     B N QA/IN 2
-----
Flags : (*) Active, (!) Administratively disabled

App : Applicant State, Bndry : Boundary,
Cls : Traffic Class, Dplx : Duplex,
Oper : MSRP Operational State, Prop : Propagated,
Reg : Registrar State
MSRP Declaration Types:
  Adv : Talker Advertise, AskFail : Listener Asking Failed,
  Fail : Talker Fail, RdyFail : Listener Ready Failed,
  Ready : Listener Ready

Applicant States:
  AA : Anxious active, AN : Anxious new,
  AO : Anxious observer, AP : Anxious passive,
  LA : Leaving active, LO : Leaving observer,
  QA : Quiet active, QO : Quiet observer,
  QP : Quiet passive, VN : Very anxious new,
  VO : Very anxious observer, VP : Very anxious passive

Registrar States:
  IN : In - Registered, LV : Leaving - Timing out,
  MT : Empty - Not Registered

#show msrp ports
21 detail Port Enabled
Oper Port Dplx Jumbo Jumbo Cls Bndry
State Sr-Pvid Speed
Size App/Reg ----
-----
----- *21 Y
Up 1000 M Full N 9216
A N QA/IN
2
2 B N QA/IN
2
Talkers: Stream Id Declaration
State Rx
Prop App Reg -----
--- 00:50:c2:4e:d3:2d:00:00 Adv Adv VO
IN 00:50:c2:4e:d3:2d:00:01 Adv Adv VO
IN
Listeners: Stream Id Declaration
State Rx
Prop App Reg -----
--- 00:50:c2:4e:d3:3d:00:00 Ready Ready VO
IN 00:50:c2:4e:d3:3d:00:01 Ready Ready VO
IN

-----
Flags : (*) Active, (!) Administratively disabled App :
Applicant State, Bndry :
: Boundary Cls : Traffic Class, Dplx :
Duplex Oper : MSRP Operational State, Prop : Propagated Reg :
Registrar State MSRP Declaration
Types: Adv :
Talker Advertise, AskFail : Listener Asking Failed, Fail : Talker
Fail, RdyFail : Listener Ready Failed, Ready : Listener
Ready Applicant
States: AA :
Anxious active, AN : Anxious new, AO :

```

```

Anxious observer,      AP      : Anxious passive,      LA      :
Leaving active,        LO      : Leaving observer,      QA      :
Quiet active,          QO      : Quiet observer,        QP      :
Quiet passive,         VN      : Very anxious new,      VO      : Very
anxious observer,     VP      : Very anxious passive    Registrar
States:      IN      : In -
Registered,    LV      : Leaving - Timing out      MT      :
Empty - Not Registered

```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on all platforms that support the AVB feature and that have an AVB feature pack license installed. To see which platforms support AVB and for information about obtaining a license, see the [Switch Engine 32.3 Feature License Requirements](#).

show msrp ports bandwidth

```
show msrp ports {port_list} bandwidth
```

Description

Displays bandwidth information of an MSRP port.

Syntax Description

msrp	Multiple Stream Registration Protocol.
ports	Ports.
<i>port_list</i>	Port list separated by a comma or "-".
bandwidth	Bandwidth information per port per traffic-class.

Default

N/A.

Usage Guidelines

Use this command to display bandwidth information of an MSRP port.

Example

```

# show msrp ports bandwidth
Port      Port      Class   Delta    Maximum    Reserved   Available
          Speed
-----
5ab       0 M   A      75.00%   0.00%      0.00%      0.00%

```

		B	0.00%	0.00%	0.00%	0.00%
*21ab	1000 M	A	75.00%	75.00%	0.00%	75.00%
		B	0.00%	75.00%	0.00%	75.00%
Flags: (*) Active, (!) Administratively disabled,						
(a) SR Class A allowed, (b) SR Class B allowed.						

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on all platforms if the AVB feature pack license is installed on the switch.

show msrp ports counters

show msrp ports {port_list} counters {event | packet}

Description

Shows PDU or event counters per port.

Syntax Description

msrp	Multiple Stream Registration Protocol.
ports	Ports.
<i>port_list</i>	Port list separated by a comma or "-".
counters	MSRP packet and attribute event counters.
event	MSRP attribute event counters.
packet	MSRP packet counters (default).

Default

N/A.

Usage Guidelines

Use this command to display PDU or event counters per port. The counters count the received attributes from talkers and listeners per attribute event, or the number of PDUs received.show msrp counters by itself displays PDU counters.

Example

#show msrp ports 17 counters packet					
Port	Streams	Reservations	Rx Pkt	Rx Error	Tx Pkt
----	-----	-----	-----	-----	-----

```
17          0          0          2          0          2

#show msrp ports 17 counters event
Port : 17
  MRP Attribute Events      Rx      Tx
  -----
  In                        250      56
  JoinIn                    0         0
  JoinMt                    224     386
  Lv                         0         0
  Mt                         0     152
  New                       0         0

  MSRP Declarations
  -----
  Listener Asking Failed    0         0
  Listener Ready            56         8
  Listener Ready Failed     0         0
  Talker Advertise          8     56
  Talker Failed             0         3
  -----

In      : Not declared, but registered
JoinIn  : Declared and Registered
JoinMt  : Declared, but not registered
Lv      : Previously registered, but now withdrawn
Mt      : Not declared, and not registered
New     : Newly declared, and possibly not previously
         registered
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on all platforms that support the AVB feature and that have an AVB feature pack license installed. To see which platforms support AVB and for information about obtaining a license, see the [Switch Engine 32.3 Feature License Requirements](#).

show msrp streams

```
show msrp streams {detail | propagation} {port port_num} {source-
  mac-addr source_mac_addr | stream-id stream_id}{destination-mac-addr
  destination_mac_addr}
```

Description

Shows the MSRP stream information collected from the talker's attributes.

Syntax Description

msrp	Multiple Stream Registration Protocol.
streams	Data streams advertising QoS specification using MSRP.
detail	Show stream information with more detail.

propagation	Show stream propagation through switch.
port	Filter based on ingress port number of the stream.
source-mac-addr	Filter based on source MAC address of a data stream.
stream-id	Filter based on stream ID of a data stream.
destination-mac-addr	Filter based on destination MAC address of a data stream.

Default

N/A.

Usage Guidelines

Use this command to show the MSRP stream information collected from the talker's attributes. The output can be filtered based on the stream id, source MAC, destination MAC, or port number on which the stream is registered.

Example

```
# show msrp streams
  Stream Id          Destination      Port  Dec  VID  Cls/Rn  BW
-----
00:50:c2:4e:d3:2d:00:00  91:e0:f0:00:35:80    17  Adv    2  A/1    6.336 Mb
00:50:c2:4e:d3:2d:00:01  91:e0:f0:00:35:81    17  Adv    2  A/1    6.336 Mb
00:50:c2:4e:d3:2d:00:02  91:e0:f0:00:35:82    17  Adv    2  A/1    6.336 Mb
00:50:c2:4e:d3:2d:00:03  91:e0:f0:00:35:83    17  Adv    2  A/1    6.336 Mb
00:50:c2:4e:d3:2d:00:04  91:e0:f0:00:35:84    17  Adv    2  A/1    6.336 Mb
Total Streams: 5

BW      : Bandwidth,           Cls      : Traffic Class,
Dec     : Prop Declaration Types, Rn       : Rank
(!)     : Talker pruned or forbidden
MSRP Declaration Types:
  Adv    : Talker Advertise,      AskFail  : Listener Asking Failed,
  Fail   : Talker Fail,          RdyFail  : Listener Ready Failed,
  Ready  : Listener Ready

#show msrp streams detail
  Stream Id          Destination      Port  Dec  VID  Cls/Rn  BW
-----
00:50:c2:4e:d3:2d:00:00  91:e0:f0:00:0e:80    17  Adv    2  A/1    6.336 Mb

      Accumulated Latency(nSec) : 0
      Max Frame Size             : 56
      Max Interval Frames        : 1
      Frame Rate (fps)           : 8000
00:50:c2:4e:d3:2d:00:01  91:e0:f0:00:0e:81    17  Fail    2  A/1    6.336 Mb
      Failure Code               : (10) Out of MSRP resrc
      Fail Bridge                : 08:00:e0:e0:e0:e0:e0:e0

      Accumulated Latency(nSec) : 0
      Max Frame Size             : 56
      Max Interval Frames        : 1
      Frame Rate (fps)           : 8000
Total Streams: 2
```

```

-----
BW      : Bandwidth,          Cls      : Traffic Class,
Dec     : Prop. Declaration Types, Rn      : Rank
(!)    : Talker pruned or forbidden

MSRP Declaration Types:
  Adv   : Talker Advertise,      AskFail  : Listener Asking Failed,
  Fail  : Talker Fail,          RdyFail  : Listener Ready Failed,
  Ready : Listener Ready

# show msrp streams propagation
  Stream Id          Destination      Port  Prop  VID  Cls/Rn      BW
-----
                                Dec
-----
00:50:c2:4e:d3:2d:00:00  91:e0:f0:00:35:80      17  Adv   2   A/1      6.336 Mb

Talker Propagation:
  Ingress    Ingress    Propagated    Propagated    Egress
  DecType    Port       DecType       Ports         DecType
  -----
  Adv        -->      17 -->  Adv        -->      19 -->  Adv
                                     21 -->  Adv

Listener Propagation:
  Egress      Egress      Propagated    Listener      Ingress
  DecType     Port       DecType       Ports         DecType
  -----
  RdyFail <--      17 <--  Ready        <--      19 <--  Ready
                                     21 <--  AskFail
                                     <--  AskFail

Total Streams: 1
-----
BW      : Bandwidth,          Cls      : Traffic Class,
Dec     : Prop. Declaration Types, Rn      : Rank
(!)    : Talker pruned or forbidden

MSRP Declaration Types:
  Adv   : Talker Advertise,      AskFail  : Listener Asking Failed,
  Fail  : Talker Fail,          RdyFail  : Listener Ready Failed,
  Ready : Listener Ready

```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on all platforms that support the AVB feature and that have an AVB feature pack license installed. To see which platforms support AVB and for information about obtaining a license, see the [Switch Engine 32.3 Feature License Requirements](#).

show msrp talkers

```

show msrp talkers {egress | ingress | ingress-and-egress} {port
  port_num}{source-mac-addr source_mac_addr | stream-id stream_id}

```


Description

Shows MSRP talker attributes.

Syntax Description

msrp	Multiple Stream Registration Protocol.
talkers	Talker attributes.
egress	Display egress talkers only (default).
ingress	Display ingress talkers only.
port	Filter based on ingress port number of the stream.
source-mac-addr	Filter based on source MAC address of a data stream.
stream-id	Filter based on stream ID of a data stream.

Default

N/A.

Usage Guidelines

Use this command to show MSRP talker attributes. The output can be filtered based on the stream id, source MAC, or port number on which the talker is registered.

Example

```
# show msrp talkers
      Stream Id          Port  Dec    Dir    State      Failure Code
-----
      App  Reg
-----
00:50:c2:4e:d3:2d:00:00  19  Adv  Egress  QA  MT  -
                        21  Fail  Egress  QA  MT  AVB incapbl port(8)
00:50:c2:4e:d3:2d:00:01  19  Adv  Egress  QA  MT  -
                        21  Fail  Egress  QA  MT  AVB incapbl port(8)
-----
App      : Applicant State,
Dir      : Direction of MSRP attribute,
Dec      : MSRP Declaration Types,
Reg      : Registrar State

MSRP Declaration Types:
  Adv    : Talker Advertise,
  Fail   : Talker Fail

Applicant States:
  AA     : Anxious active,
  AO     : Anxious observer,
  LA     : Leaving active,
  QA     : Quiet active,
  QP     : Quiet passive,
  VO     : Very anxious observer,
  AN     : Anxious new,
  AP     : Anxious passive,
  LO     : Leaving observer,
  QO     : Quiet observer,
  VN     : Very anxious new,
  VP     : Very anxious passive

Registrar States:
  IN     : In - Registered,
  MT     : Empty - Not Registered
  LV     : Leaving - Timing out,
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on all platforms that support the AVB feature and that have an AVB feature pack license installed. To see which platforms support AVB and for information about obtaining a license, see the [Switch Engine 32.3 Feature License Requirements](#).

show mvr

```
show mvr {vlan vlan_name}
```

Description

Displays the MVR configuration on the switch.

Syntax Description

<i>vlan_name</i>	Specifies a <u>VLAN</u> name.
------------------	-------------------------------

Default

N/A.

Usage Guidelines

If a VLAN is specified, information for the VLAN is displayed.

Example

The following command displays the MVR configuration for the VLAN accounting:

```
show mvr accounting
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MVR feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

show mvr cache

This command is provided for backward compatibility. The recommended command is:

```
show mcast cache {{vlan} vlan_name} {[group grpaddressMask |
grpaddressMask] {source sourceIP | sourceIP}} {type [snooping | pim
| mvr]}| {summary}}
```

The syntax for the original form of this command is:

```
show mvr cache {vlan vlan_name}
```

Description

Displays the multicast cache entries added by MVR.

Syntax Description

vlan_name	Specifies a <u>VLAN</u> name.
-----------	-------------------------------

Default

N/A.

Usage Guidelines

If no VLAN is specified, information for all VLANs is displayed.

Example

The following command displays the multicast cache in the MVR range for the VLAN vlan110:

```
Switch.78 # show mvr cache vlan110
```

This command display is the same as for the following preferred command:

```
show mcast cache {{vlan} name} {[group grpaddressMask | grpaddressMask]
{source sourceIP | sourceIP}} {type [snooping | pim | mvr]}| {summary}}
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MVR feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

show mvrp

show mvrp

Description

This command is used to show the mvrp settings as follows. If the MVRP enabled port is a load shared port, a suffix 'g' is displayed.

Syntax Description

mvrp	Multiple <u>VLAN</u> Registration Protocol.
-------------	---

Default

N/A.

Usage Guidelines

Use this command to show the MVRP settings.

Example

```
# show mvrp
MVRP enabled           : Enabled
MVRP dynamic VLAN creation : Enabled
MVRP VLAN registration  : Forbidden
MVRP default STP domain   : s0
MVRP enabled ports      : 9      *11    *13
Flags: (*) Active, (!) Administratively disabled.
      (g) Load Sharing Port
      (G) Multi-switch LAG Group port
```

History

This command was first available in ExtremeXOS 15.3.

MRVP VLAN registration output was added in 15.3.2.

Flag showing if the MVRP-enabled port is an port added in ExtremeXOS 22.1

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show mvrp ports counters

show mvrp ports {port_list} **counters** {event | packet}

Description

Shows the port MVRP statistics. The statistics for packet or event counters are displayed as per input.

Syntax Description

mvrp	Multiple <u>VLAN</u> Registration Protocol.
ports	Ports.
<i>port_list</i>	List of ports in the switch separated by a comma or "-".
counters	MVRP packet and attribute event counters.
event	MVRP attribute event counters.
packet	MVRP packet counters (default).

Default

Packet counters.

Usage Guidelines

This command is to show the port MVRP statistics. The statistics for packet or event counters will be displayed as per input. The default is packet counters. The packet counters include Number of VLANs registered on the port, Number of Failed Registrations, Number of MVRPDUs received, Number of MVRPDUs sent, Number of erroneous MVRPDUs received, and the source address of the MVRP message last received by the port. The event counters include the number of different events received/transmitted.

Example

```
# show mvrp ports 9,11,13 counters packet
Port      VLANs   Failed   Rx Pkt  Rx Error  Tx Pkt   Last Source
      Regs   Regs   Count   Count    Count   Address
-----
   9       2       0       0        0        64   00:00:00:00:00:00
  11       2       0   806836        0   433754   00:22:97:00:41:e7
  13       2       0   784176        0   404794   00:22:97:00:41:e8

-----
Regs: Registrations

# show mvrp ports 9 counters event
Port : 17
      MRP Attribute Events      Rx      Tx
-----
In                250        56
JoinIn             0         0
JoinMt            224       386
LeaveAll            5         0
Lv                 0         0
Mt                 0       152
New

-----
In      : Not declared, but registered
```

```

JoinIn   : Declared and Registered
JoinMt   : Declared, but not registered
LeaveAll  : All registrations will shortly be deregistered
Lv       : Previously registered, but now withdrawn
Mt       : Not declared, and not registered
New      : Newly declared, and possibly not previously registered

```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show mvrp tag

```
show mvrp tag vlan_tag {ports {port_list}}
```

Description

Shows the port specific applicant and registrar states and the configured control values for all MVRP enabled ports.

Syntax Description

mvrp	Multiple <u>VLAN</u> Registration Protocol.
tag	The 802.1Q VLAN ID.
<i>vlan_tag</i>	VLAN ID ranging from 1 to 4094 (default is 2).
<i>port_list</i>	Port list separated by comma or "-".

Default

N/A.

Usage Guidelines

Use this command to show the port specific applicant and registrar states and the configured control values for all MVRP enabled ports. The registrar control value is derived as follows:

- Normal = Dynamically ordered port.
- Fixed = Statically added port.
- Forbidden = VLAN is configured to be forbidden on the port.

Example

```
# show mvrp tag 2
```

Port	Applicant State	Applicant Control	Registrar State	Registrar Control	
	-----	-----	-----	-----	-----
	9	VN	On	MT	Normal
	11	QA	On	IN	Normal
	13	QA	On	IN	Normal
Applicant States:					
	AA	: Anxious active,	AN	: Anxious new,	
	AO	: Anxious observer,	AP	: Anxious passive,	
	LA	: Leaving active,	LO	: Leaving observer,	
	QA	: Quiet active,	QO	: Quiet observer,	
	QP	: Quiet passive,	VN	: Very anxious new,	
	VO	: Very anxious observer,	VP	: Very anxious passive	
Registrar States:					
	IN	: In - Registered,	LV	: Leaving - Timing out,	
	MT	: Empty - Not Registered			
Applicant Control:					
	On	: Transmit On,	Off	: Transmit Off	
Registrant Control:					
	Fixed	: Statically added,	Forbidden	: Forbidden VLAN,	
	Normal	: Dynamically added			

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show netlogin

```
show netlogin {port port_list [ {vlan} vlan_name | vlan vlan_list]}
               {dot1x {detail}} {mac} {web-based}
```

Description

Shows status information for network login.

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
<i>vlan_name</i>	Specifies the name of a <i>VLAN</i> .
<i>vlan_list</i>	Specifies a VLAN list of IDs.
dot1x	Specifies 802.1X information.
detail	Shows detailed information.
mac	Specifies MAC-based information.
web-based	Specifies web-based information.

Default

N/A.

Usage Guidelines

Depending on your configuration, software version, and the parameters you choose to display, the information reported by this command may include some or all of the following:

- Whether network login is enabled or disabled.
- The base-URL.
- The default redirect page.
- The logout privileges setting.
- The network login session-refresh setting and time.
- The MAC and IP address of supplicants.
- The type of authentication, 802.1X, MAC-based, or HTTP (web-based).
- The guest VLAN configurations, if applicable.
- The dynamic VLAN state and uplink ports, if configured.
- Whether network login port restart is enabled or disabled.
- Which order of authentication protocols is currently being used.

If you do not specify the authentication method, the switch displays information for all network login authentication methods.



Note

The "current user" count displays how many resources are left to be able to be configured/authenticated. Admin-profile rules consume a resource similar to authenticated users, even if that particular MAC address is not presently on the system (a static admin-profile port rule also increments this count). As a result, the "current user" count value reflects a combination of users and consumed resources (admin-profile rules).

Example

The following sample output shows the summary network login information:

```
# show netlogin

NetLogin Authentication Mode : web-based ENABLED; 802.1X ENABLED; mac-based ENABLED
NetLogin VLAN                : "nvlan"
NetLogin move-fail-action    : Authenticate
NetLogin Client Aging Time   : 5 minutes
Dynamic VLAN Creation        : Enabled
Dynamic VLAN Uplink Ports    : 12
-----
Web-based Mode Global Configuration
-----
Base-URL                     : network-access.com
Default-Redirect-Page        : http://www.yahoo.com
Logout-privilege              : YES
Netlogin Session-Refresh     : ENABLED; 3 minutes
Authentication Database       : Radius, Local-User database
-----
-----
```


802.1X Mode Global Configuration

```

-----
Quiet Period                : 60
Supplicant Response Timeout : 30
Re-authentication period    : 200
RADIUS server timeout       : 30
EAPOL MPDU version to transmit : v1
Authentication Database      : Radius
-----

```

MAC Mode Global Configuration

```

-----
Re-authentication period    : 0 (Re-authentication disabled)
Authentication Database      : Radius, Local-User database
Authentication Delay Period : 0 (Default)
-----

```

MAC Address/Mask	Password (encrypted)	Port(s)
00:00:86:3F:1C:35/48	yaqu	any
00:01:20:00:00:00/24	yaqu	any
00:04:0D:28:45:CA/48	=4253C5;500@	any
00:10:14:00:00:00/24	yaqu	any
00:10:A4:A9:11:3B/48	yaqu	any
00:10:A4:00:00:00/24	yaqu	any
Default	yaqu	any

Authentication Database : Radius, Local-User database

Port: 5, Vlan: nvlan, State: Enabled, Authentication: mac-based, Guest Vlan <Not Configured>: Disabled

MAC	IP address	Authenticated	Type	ReAuth-Timer	User

Port: 9, Vlan: nvlan, State: Enabled, Authentication: web-based, Guest Vlan <Not Configured>: Disabled

MAC	IP address	Authenticated	Type	ReAuth-Timer	User

Port: 10, Vlan: nvlan, State: Enabled, Authentication: 802.1X, mac-based, Guest Vlan <Not Configured>: Disabled

MAC	IP address	Authenticated	Type	ReAuth-Timer	User

Port: 17, Vlan: engr, State: Enabled, Authentication: mac-based, Guest Vlan <Not Configured>: Disabled

MAC	IP address	Authenticated	Type	ReAuth-Timer	User

Port: 17, Vlan: mktg, State: Enabled, Authentication: mac-based, Guest Vlan <Not Configured>: Disabled

MAC	IP address	Authenticated	Type	ReAuth-Timer	User

Port: 19, Vlan: corp, State: Enabled, Authentication: 802.1X, Guest Vlan <Not Configured>: Disabled

MAC	IP address	Authenticated	Type	ReAuth-Timer	User
00:04:0d:50:e1:3a	0.0.0.0	No		0	00040D50E13A
00:10:dc:98:54:00	10.201.31.113	Yes, Radius	802.1X	24	md5isp7

Port: 19, Vlan: nvlan, State: Enabled, Authentication: 802.1X, Guest Vlan <Not Configured>: Disabled

MAC	IP address	Authenticated	Type	ReAuth-Timer	User
00:04:0d:50:e1:3a	0.0.0.0	No	802.1X	0	

Port: 19, Vlan: voice-ip, State: Enabled, Authentication: 802.1X, Guest Vlan <Not Configured>: Disabled

MAC	IP address	Authenticated	Type	ReAuth-Timer	User
00:04:0d:50:e1:3a	0.0.0.0	Yes, Radius	802.1X	75	00040D50E13A

The following command shows more detailed information, including the configured authentication methods:

```
# show netlogin port 3:2 vlan "Default"
Port: 2:1          Vlan: Default
Authentication: Web-Based, 802.1X
Port State:       Unauthenticated
Guest VLAN:       Not Enabled
DHCP:             Not Enabled
MAC               IP address      Auth   Type       ReAuth-Timer User
00:0C:F1:E8:4E:13 0.0.0.0         No    802.1X     0          Unknown
00:01:30:F3:EA:A0 10.0.0.1        Yes   802.1X     0          testUser
```

The following command shows information about a specific port configured for network login:

```
# show netlogin port 1:1
Port      : 1:1
Port Restart : Enabled
Vlan      : Default
Authentication: mac-based
Port State : Enabled
Guest Vlan : Disabled
MAC       IP address      Auth   Type       ReAuth-Timer   User
-----
```

The following command shows information for 802.1X mode:

```
# show netlogin dot1x
NetLogin Authentication Mode : web-based DISABLED; 802.1x ENABLED; MAC-based ENABLED
NetLogin VLAN                : "nlvlan"
NetLogin move-fail-action    : Deny
NetLogin Client Aging Time   : 5 minutes
Dynamic VLAN Creation        : Disabled
Dynamic VLAN Uplink Ports    : None
Authentication Protocol Order: 802.1x, web-based, mac-based (default)
Maximum Number Of Users      : 256 (Policy Enabled only)

-----
                        802.1x Mode Global Configuration
-----
EAPOL MPDU version to transmit : v1
Tag EAPOL on tagged ports      : Off
Authentication Database        : Radius
RADIUS Accounting              : On
-----

Port: 1, Vlan: nlvlan, State: Enabled, Authentication: 802.1x, mac-based
Authentication Failure Vlan <Not Configured>: Disabled
Authentication Service-Unavailable Vlan <Not Configured>: Disabled
-----
                        802.1x Port Configuration
-----
Quiet Period                   : 60
Supplicant Response Timeout    : 30
Re-authentication              : On
Re-authentication period       : 1200
Max Re-authentications         : 3
RADIUS server timeout          : 30
Guest Vlan <Not Configured>    : Disabled
-----
                        MAC Mode Port Configuration
-----
Re-authentication              : Off
Re-authentication period       : 3600
```

```

Authentication Delay      : 0 seconds (Default)
-----
Netlogin Clients
-----
MAC                        IP address      Authenticated   Type    ReAuth-Timer  User
-----
(B) - Client entry Blackholed in FDB

Port: 1, Vlan: v1, State: Enabled, Authentication: 802.1x, mac-based
Authentication Failure Vlan <Not Configured>: Disabled
Authentication Service-Unavailable Vlan <Not Configured>: Disabled
-----
802.1x Port Configuration
-----
Quiet Period              : 60
Supplicant Response Timeout : 30
Re-authentication         : On
Re-authentication period   : 1200
Max Re-authentications     : 3
RADIUS server timeout      : 30
Guest Vlan <Not Configured>: Disabled
-----
MAC Mode Port Configuration
-----
Re-authentication         : Off
Re-authentication period   : 3600
Authentication Delay      : 0 seconds (Default)
-----
Netlogin Clients
-----
MAC                        IP address      Authenticated   Type    ReAuth-Timer  User
-----
00:00:00:00:00:02 0.0.0.0        Yes, Radius     802.1x  658           harish
-----
(B) - Client entry Blackholed in FDB

Number of Clients Authenticated : 1

```

For 802.1X, if re-authentication is disabled, the re-authentication period appears as follows:

```

Re-authentication period : 0 (Re-authentication disabled)

```

The `show netlogin port 5:4 dot1x` command generates the following sample output:

```

Port                : 5:4
Port Restart        : Disabled
Vlan                : corp
Authentication      : 802.1X
Port State          : Enabled
Guest Vlan          : Enabled
MACIP addressAuthenticatedTypeReAuth-TimerUser
00:10:dc:92:53:2d10.201.31.119Yes,Radius802.1X14md5isp4
-----

```

The `show netlogin port 5:4 dot1x detail` command generates the following sample output:

```

Port: 5:4
Port Restart: Disabled
Vlan: corp
Authentication: 802.1X

```

```

Port State: Enabled
Guest Vlan: Enabled
MAC
00:10:dc:92:53:2d : IP=10.201.31.119 Auth=Yes User=md5isp4
: AuthPAE state=AUTHENTICATED BackAuth state=IDLE
: ReAuth time left=8 ReAuth count=0
: Quiet time left=0
-----

```

History

This command was first available in ExtremeXOS 11.1.

Information about the guest VLAN was added in ExtremeXOS 11.2.

Information about the configured port MAC list was added in ExtremeXOS 11.3.

Information about dynamic VLANs and network login port restart was added in ExtremeXOS 11.6.

The `vlan_list` variable was added in ExtremeXOS 16.1.

Information about authentication delay added in ExtremeXOS 21.1.

Authentication username format information was added in ExtremeXOS 22.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show neighbor-discovery cache ipv6

```

show neighbor-discovery {cache {ipv6}} {[ipv6_addr | mac | permanent]
    {vr vr_name}} | [ {vlan vlan_name | vlan vlan_list] | vr vr_name}

```

Description

This command displays all the entries from the neighbor cache.

Syntax Description

<i>ipv6_addr</i>	Specifies an IPv6 address.
<i>mac</i>	Specifies a MAC address.
permanent	Specifies static entries.
<i>vr_name</i>	Specifies a VR or VRF.
<i>vlan_name</i>	Specifies an IPv6 configured <u>VLAN</u> .
<i>vlan_list</i>	Specifies an IPv6 configured VLAN list of IDs.

Default

N/A.

Usage Guidelines

This command displays the entries present in the neighbor cache.

The entries displayed can be filtered by IPv6 address, MAC address, or by VLAN. The permanent keyword filters the output to display static entries.

The `vr_name` indicates the VR or VRF on which the operation is performed. In its absence, the operation applies to VR-Default.

Example

The following example output shows all entries from the neighbor cache:

```
# show neighbor-discovery cache ipv6
VR          Destination
Mac         Age  Static  VLAN      VID    Port
VR-Default  2001:db8:100::7
00:01:30:00:6b:00    0      NO   gtag100    100    1:2
VR-Default  2001:db8:100::99
00:01:02:33:33:33    0      YES  gtag100    100
VR-Default  2001:db8:99::99
00:01:02:01:01:01    0      YES  gtag99     99
Total Entries      :      0
Dynamic Entries    :      0      Static Entries      :      0
Pending Entries    :      0
Timeout           :      20 minutes  Refresh           :      Enable
Neighbor Discovery Global Settings
Max Entries        :      4096
Max Pending Entries :      1024
```

History

This command was first available in ExtremeXOS 11.2.

The `vlan_list` variable was added in ExtremeXOS 16.1.

Max entries and max pending entries moved under a global settings heading in ExtremeXOS 30.1.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

show netlogin authentication failure vlan

```
show netlogin authentication failure [ {vlan} vlan_name | vlan
vlan_list]
```

Description

Displays the authentication failure VLAN related configuration details.

Syntax Description

<i>vlan_name</i>	Specifies the name of a failure VLAN.
<i>vlan_list</i>	Specifies the failure VLAN list of IDs.

Default

N/A.

Usage Guidelines

Use this command to display configuration details for the authentication failure VLAN.

Example

The following is sample output from this command:

```
# show netlogin authentication failure vlan

-----
Authentication Service unavailable
Vlan
port                Status
-----
corp
1:2                 Disabled
```

History

This command was first available in ExtremeXOS 12.1.

The *vlan_list* variable was added in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show netlogin authentication service-unavailable vlan

```
show netlogin authentication service-unavailable [ {vlan} vlan_name |
vlan vlan_list]
```

Description

Displays the authentication service-unavailable VLAN related configuration details.

Syntax Description

<i>vlan_name</i>	Specifies the name of the authentication service-unavailable VLAN.
<i>vlan_list</i>	Specifies the authentication service-unavailable VLAN list of IDs.

Default

N/A.

Usage Guidelines

Use this command to display configuration details for the service-unavailable VLAN.

Example

The following is sample output from this command:

```
# show netlogin authentication service-unavailable vlan serv
```

```
-----  
-  
Authentication Service unavailable VLAN      Port      Status  
-----  
-  
serv                                         1 (untagged)  Enabled
```

History

This command was first available in ExtremeXOS 12.1.

The *vlan_list* variable was added in ExtremeXOS 16.1.

Additional information associated with the ability to configure multiple service-unavailable VLANs added in ExtremeXOS 30.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show netlogin banner

```
show netlogin banner
```

Description

Displays the user-configured banner string for network login.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to view the banner that is displayed on the network login page.

Example

The following command displays the network login banner:

```
show netlogin banner
```

If a custom banner web page exists, show banner netlogin generates the following output:

```
***** Testing NETLOGIN BANNER at <system name>*****
NOTE: Banner is not in use. Overridden since custom login page "netlogin_login_page.html"
is present.
```

If a custom banner web page does not exist, nothing is displayed.

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show netlogin guest-vlan

```
show netlogin guest-vlan [ {vlan} vlan_name | vlan vlan_list]
```

Description

Displays the configuration for the guest VLAN feature.

Syntax Description

<i>vlan_name</i>	Specifies the name of a guest VLAN.
<i>vlan_list</i>	Specifies the guest VLAN list of IDs.

Default

N/A.

Usage Guidelines

Use this command to display the guest VLANs configured on the switch.

If you specify the `vlan_name`, the switch displays information for only that guest VLAN.

The output displays the following information in a tabular format:

- Port—Specifies the 802.1X enabled port configured for the guest VLAN.
- Guest-vlan—Displays the enabled/disabled state of the guest VLAN feature.
- Vlan—Specifies the name of the guest VLAN.

Example

The following sample output displays the local network login list:

```
# show netlogin guest-vlan
```

Port	Guest-vlan	Vlan
5:1	Disabled	gv11
5:2	Enabled	gv12
5:3	Disabled	gv13
5:4	Enabled	gv14

History

This command was first available in ExtremeXOS 11.6.

The `vlan_list` variable was added in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show netlogin local-users

```
show netlogin local-users
```

Description

Displays the local network login users configured on the switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to display the list of local network login users and associated VLANs.

If you associated a VLAN with a local network login user, the output displays the name of the VLAN. If you have not associated a VLAN with a local network login user, the output displays not configured.

The Extended-VLAN VSA column displays the name of the VLAN and the following information:

- *not configured*—Specifies that you have not associated a VLAN with a local network login user.
- *—Specifies the movement based on the incoming port's traffic. For example, the VLAN behaves like VSA 203 if identified with a VLAN name or VSA 209 if identified with a VLAN ID.
- T—Specifies a tagged client.
- U—Specifies an untagged client.

In addition, this output is useful to determine which local network login user you want to modify or delete from the system.

Example

The following command displays the local network login list:

```
show netlogin local-users
```

The following is sample output from this command:

Netlogin Local User Name	Password (encrypted)	Extended-VLAN VSA
-----	-----	-----
000000000012	Iqyydz\$MP7AG.VAmwOoqiKX2u13H1	U hallo
00008653C314	Bo028L\$orVvKv8.wmxcorhhXxQY40	* default
megtest	w7iMbp\$lBL34/dLx4G4M8aAdiCvI	<not configured>
testUser	/Jhouw\$iHEl5steebwhOibgj6pZq.	T testVlan

History

This command was first available in ExtremeXOS 11.2.

The output was modified to include VLAN information in ExtremeXOS 11.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show netlogin mac-list

```
show netlogin mac-list
```

Description

Displays the MAC address list for MAC-based network login.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to display the MAC address list used for MAC-based network login.

MAC-based authentication is VR aware, so there is one MAC list per VR.

Example

The following command displays the MAC address list:

```
show netlogin mac-list
```

The following is sample output from this command:

MAC Address/Mask	Password (encrypted)	Port(s)
00:00:00:00:00:10/48	<not configured>	1:1-1:5
00:00:00:00:00:11/48	<not configured>	1:6-1:10
00:00:00:00:00:12/48	<not configured>	any
00:01:30:70:0C:00/48	yaquany	
00:01:30:32:7D:00/48	ravdqsrany	
00:04:96:00:00:00/24	<not configured>	any

History

This command was first available in ExtremeXOS 11.1.

Information about the configured port MAC list was added in ExtremeXOS 11.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show netlogin session

```
show netlogin session {all | summary} {mac-address mac_address} {ports
ports} {agent [convergence-endpointdot1x | mac | web-based]}
```

Description

Use this command to display the active authenticated sessions for both policy and non-policy mode.

Syntax Description

all	Include sessions with 'terminated' status.
summary	Do not display detailed information.
mac-address	Specify a MAC address.
<i>mac_address</i>	Specify a MAC address.
ports	Port or range of ports.
<i>ports</i>	Port or range of ports.
agent	Specify an agent type.
dot1x	IEEE 802.1X Port-Based Network Access Control.
mac	MAC authentication.
web-based	Web-based authentication.
convergence-endpoint	Convergence-endpoint authentication.

Default

N/A.

Usage Guidelines

None.

Example

The following example shows a summary of the active authenticated sessions:

```
# show netlogin session summary
Port      Address type Address                               X W M C B
-----
23        mac      00:00:88:88:00:00                      X
23        mac      00:00:88:88:00:01                      X  X
23        mac      00:00:88:88:00:02                      X  X
23        mac      00:00:88:88:00:03                      X  X
23        mac      00:00:88:88:00:04                      X
-----
X - dot1x
W - web-based
M - mac
C - convergence-endpoint
B - Blackholed in FDB
```

The following example shows information about the authentication sessions:

```
show netlogin session
Multiple authentication session entries
-----
Port      : 1:1      Station address : 00:04:f2:9f:48:fe
Auth status : success Last attempt   : Fri Nov 17 21:56:01 2017
Agent type  : cep    Session applied : true
Server type : local  VLAN-Tunnel-Attr : None
Policy index : 2     Policy name    : NProfile_16 (active)
Session timeout : 0   Session duration : 2 days, 21:57:19
```

```

Idle timeout      : 300      Idle time       : 0:00:00
Auth-Override     : enabled  Termination time : Not Terminated

Port              : 15       Station address  : 00:00:00:00:0b:00
Auth status       : success  Last attempt     : Thu Jan 30 12:48:04 2020
Agent type        : mac      Session applied  : false
Server type       : local    VLAN-Tunnel-Attr : None
Policy index      : 0        Policy name      : No Policy applied
Session timeout   : 0        Session duration : 0:00:04
Idle timeout      : 0        Idle time       : 0:00:00
Auth-Override     : disabled Termination time : Not Terminated

```

History

This command was first available in ExtremeXOS 16.1.

Sessions shown for both policy and non-policy mode was added in ExtremeXOS 30.7.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show netlogin timeout

```
show netlogin timeout
```

Description

Use this command to display the *NetLogin* timeout value for an idle or session timeout.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

This example displays the `show netlogin timeout` command output.

```

#show netlogin timeout
      Session      Idle
Authentication  Timeout  Timeout
Type            (sec)   (sec)
-----
dot1x           0       300

```

web-based	0	300	
mac	0	300	
convergence-endpoint		0	300

History

This command was first available in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show netlogin trap

```
show netlogin trap
```

Description

Use this command to display *NetLogin* trap settings.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

```
# show netlogin trap

System: Disabled

Enabled-Trap Ports
Success      :
Failed       :
Terminated   :
Max-Reached :

Disabled-Trap Ports
Success      : 1-3
Failed       : 1-3
Terminated   : 1-3
Max-Reached : 1-3
```

History

This command was first available in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show network-clock gntp

show network-clock gntp {**default-set** | **current-set** | **parent-set** | **time-properties-set**}

Description

Displays global gPTP configuration and data.

Syntax Description

default-set	Displays this switch's native time capabilities.
current-set	Displays this switch's state relative to the grandmaster system.
parent-set	Displays the upstream (i.e., toward the grandmaster) system's parameters.
time-properties-set	Displays the grandmaster's parameters.

Default

N/A.

Usage Guidelines

Use this command to display whether gPTP is enabled in the switch and the ports on which gPTP is enabled.

Example

```
# show network-clock gntp
gPTP status      : Enabled
BMCA             : Off
Static slave port : None
gPTP enabled ports : *1m      *21d      *22d      *47d

Flags:           (*) Active, (!) Administratively disabled,
                  (d) Disabled gPTP port role, (m) Master gPTP port role,
                  (p) Passive gPTP port role,  (s) Slave gPTP port role

# show network-clock gntp default-set
Local Clock Identity      : 00:04:96:FF:FE:52:2C:BE
Number of gPTP ports     : 24
```

```

Local Clock Class           : 255 (slave only clock)
Local Clock Accuracy        : 254 (unknown)
Local Offset Scaled Log Variance : 65535
GM Capable                  : No
Local Priority1              : 255
Local Priority2              : 248
Current UTC Offset          : unknown
Leap 59                     : No
Leap 61                     : No
Time Traceable              : No
Frequency Traceable         : No
Time Source                  : 160 (Internal Oscillator)

# show network-clock gtp current-set
Steps Removed               : 1
Offset from GM              : 10 nanoseconds
Last GM Phase Change        : 548 nanoseconds
Last GM Frequency Change    : 100
GM Time Base Indicator      : 2
GM Change Count             : 1
Last GM Change Event        : Tue Nov 22 03:32:07 2011
Last GM Frequency Change Event : Tue Nov 22 03:32:07 2011
Last GM Phase Change Event  : Tue Nov 22 03:32:07 2011

# show network-clock gtp parent-set
Parent Clock Identity        : 00:04:96:FF:FE:52:34:5F
Parent port number          : 21
Cumulative Rate Ratio        : 10000
GM Clock Identity           : 00:12:34:FF:FE:56:78:9A
GM Clock Accuracy            : 32 (25 ns)
GM Offset Scaled Log Variance : 32767
GM Priority1                 : 245
GM Priority2                 : 248

# show network-clock gtp time-properties-set
Current UTC Offset          : 33 seconds
Leap 59                     : No
Leap 61                     : No
Time Traceable              : Yes
Frequency Traceable         : Yes
Time Source                  : 32 (GPS)

```

History

This command was first available in ExtremeXOS 15.3.

Support for BCMA was first available in ExtremeXOS 15.7.1.

Platform Availability

This command is available on all platforms that support the AVB feature and that have an AVB feature pack license installed. To see which platforms support AVB and for information about obtaining a license, see the [Switch Engine 32.3 Feature License Requirements](#).

show network-clock gtp ports

```
show network-clock gtp ports [port_list | all] {counters}
```


Description

Displays gPTP port parameters and counters.

Syntax Description

<i>port_list</i>	Specifies one or more of the switch's physical ports.
all	Specifies all of the switch's physical ports.

Default

N/A.

Usage Guidelines

The command `show network-clock gtp port` displays the specified port's gPTP parameters:

Physical port number	The switch's number for this physical port.
gPTP port status	Indicates whether gPTP is enabled on this port.
Clock Identity	This switch's gPTP Clock Identity.
gPTP Port Number	gPTP number for this physical port.
IEEE 802.1AS Capable	Indicates whether this switch and the neighboring system device connected via this port can interoperate via gPTP.
Port Role	The port's gPTP role: <ul style="list-style-type: none"> • Disabled (3) • Master (6) • Passive (7) • Slave (9)
Announce Initial Interval	The initial announce interval on this port. The interval is represented as the log base 2 of the interval in seconds; for example, 0 = 1 second.
Announce Current Interval	The current announce interval on this port. The interval is represented as the log base 2 of the interval in seconds; for example, 1 = 2 seconds.
Announce Receipt Timeout	The number of announce intervals a slave port waits without receiving an Announce message before it assumes the master port is no longer sending Announce messages and the BMCA needs to be run.
Sync Initial Interval	The initial time-synchronization transmission interval on this port. The interval is represented as the log base 2 of the interval in seconds; for example, -1 = 500 milliseconds.
Sync Current Interval	The current time-synchronization transmission sync interval on this port. The interval is represented as the log base 2 of the interval in seconds; for example, -2 = 250 milliseconds.
Sync Receipt Timeout	The number of time-synchronization transmission intervals a slave port waits without receiving a Sync message before it assumes the master port is no longer sending Sync messages and the BMCA needs to be run.
Sync Receipt Timeout Interval	Sync Receipt Timeout in time units.

Peer Delay Initial Interval	The initial Peer Delay Request interval on this port. The interval is represented as the log base 2 of the interval in seconds; for example, 2 = 4 seconds.
Peer Delay Current Interval	The current Peer Delay Request interval on this port. The interval is represented as the log base 2 of the interval in seconds; for example, 3 = 8 seconds.
Peer Delay Allowed Lost Responses	The number of consecutive Peer Delay Request messages that the switch must send on this port without receiving a valid response before it considers the port not to be exchanging Peer Delay messages with its neighbor.
Measuring Propagation Delay	Indicates whether this port is measuring its link's propagation delay.
Mean Propagation Delay	The link's estimated one-way propagation delay. The peer delay protocol measures the sum of the link's propagation delays in each direction, and this is that sum divided by two, which is accurate only if the link is symmetrical.
Mean Propagation Delay Threshold	The propagation delay above which the switch considers this port unable to run gPTP.
Propagation Delay Asymmetry	The configured time that the propagation delay from this switch to the neighbor is less than the estimated one-way propagation delay between the switch and its neighbor (which is also the time that the propagation delay from the neighbor to this switch is greater than the estimate). This value is negative if the propagation delay to the neighbor is greater than the estimate. Let tIR be the propagation delay from this switch (initiator) to the neighbor (responder), tRI be the propagation delay from the neighbor to this switch, and meanPathDelay be the estimated one-way propagation delay. Then: $\text{meanPathDelay} = (\text{tIR} + \text{tRI}) / 2$ $\text{tIR} = \text{meanPathDelay} - \text{asymmetry_time}$ $\text{tRI} = \text{meanPathDelay} + \text{asymmetry_time}$
Neighbor Rate Ratio	The estimated ratio of the frequency of the local clock in the neighboring systemdevice connected via this port, to this switch's local clock's frequency. The ratio is represented as the ratio minus 1, multiplied by 2^{41} : $(\text{ratio} - 1) * 2^{41}$
PTP Version	The PTP version number used on this port. Always 2.
Peer Delay Correction Field Fractional nano seconds only	Shows whether or not you consider only the fractional nano-second portion of correction field of peer delay messages. Default is off.

The command `show network-clock gptp port counters` displays the specified port's gPTP counters:

Physical port number	The switch's number for this physical port.
gPTP port status	Indicates whether gPTP is enabled on this port.
Announce	The number of Announce messages received and sent.
Sync	The number of Sync messages received and sent.
Follow Up	The number of Follow Up messages received and sent.
Peer Delay Request	The number of Peer Delay Request messages received and sent.
Peer Delay Response	The number of Peer Delay Response messages received and sent.

Peer Delay Response Followup	The number of Peer Delay Response Follow Up messages received and sent
gPTP packet discards	The number of received gPTP packets discarded or lost for one of the following reasons (from 802.1AS-2011 14.7.8): <ul style="list-style-type: none"> • Announce message from this switch • Announce message with stepsRemoved >= 255 • Announce message with a Path Trace TLV that includes this switch • Follow Up message not received following Sync message received • Peer Delay Response message not received following Peer Delay Request message sent • Peer Delay Response Follow Up message not received following Peer Delay Request message sent
Announce Receipt Timeout Count	The number of Announce Receipt timeouts.
Sync Receipt Timeout Count	The number of Sync Receipt timeouts.
Peer Delay Allowed Lost Responses Exceeded Count	The number of times the number of consecutive Peer Delay Request messages sent without receiving a valid response exceeded the Peer Delay Allowed Lost Responses.

Example

```
# show network-clock gptp ports 2
Physical port number      : 2
gPTP port status         : Enabled
Clock Identity           : 00:04:96:FF:FE:52:2C:BE
gPTP Port Number        : 2
IEEE 802.1AS Capable     : Yes
Port Role                : 9 (Slave)
Announce Initial Interval : 0 (1 second)
Announce Current Interval : 1 (2 seconds)
Announce Receipt Timeout  : 3
Sync Initial Interval     : -3 (125 milliseconds)
Sync Current Interval     : -2 (250 milliseconds)
Sync Receipt Timeout      : 3
Sync Receipt Timeout Interval : 750 milliseconds
Peer Delay Initial Interval : 2 (4 seconds)
Peer Delay Current Interval : 4 (8 seconds)
Peer Delay Allowed Lost Responses : 3
Measuring Propagation Delay : Yes
Mean Propagation Delay    : 1000 nanoseconds
Mean Propagation Delay Threshold : 10000 nanoseconds
Propagation Delay Asymmetry : 0
Neighbor Rate Ratio       : 200
PTP Version               : 2
Peer Delay Correction Field
  Fractional nano seconds only : On

# show network-clock gptp ports 3 counters
Physical port number      : 3
gPTP port status         : Enabled
-----
Parameter                Receive      Transmit
-----
Announce                  1000        2000
Sync                      1000        500
```

Follow Up	2000	2500
Peer Delay Request	3000	1000
Peer Delay Response	500	1500
Peer Delay Response Follow Up	200	1000
gPTP packet discards	2000	-

Announce Receipt Timeout Count		: 1000
Sync Receipt Timeout Count		: 500
Peer Delay Allowed Lost Responses Exceeded Count		: 2000

History

This command was first available in ExtremeXOS 15.3.

Whether or not you consider only the fractional nano-second portion of correction field of peer delay messages information was added in ExtremeXOS 31.1.

Platform Availability

This command is available on all platforms that support the AVB feature and that have an AVB feature pack license installed. To see which platforms support AVB and for information about obtaining a license, see the [Switch Engine 32.3 Feature License Requirements](#).

show node

show node {**detail**}

Description

Displays the status of the nodes in the system as well as the general health of the system.

Syntax Description

detail	Displays the information on a per-node basis rather than in a tabular format.
---------------	---

Default

N/A.

Usage Guidelines

Use this command to display the current status of the nodes and the health of the system. The information displayed shows the node configurations (such as node priority) and the system and hardware health computations. You can use this information to determine which node will be elected primary in case of a failover.

The following table lists the node statistic information collected by the switch.

Table 40: Node States

Node State	Description
BACKUP	In the backup state, this node becomes the primary node if the primary fails or enters the DOWN state. The backup node also receives the checkpoint state data from the primary.
DOWN	In the down state, the node is not available to participate in leader election. The node enters this state during any user action, other than a failure, that makes the node unavailable for management. Examples of user actions are: <ul style="list-style-type: none"> Upgrading the software Rebooting the system using the <code>reboot</code> command. Synchronizing the software and configuration in non-volatile storage using the <code>synchronize</code> command
FAIL	In the fail state, the node has failed and needs to be restarted or repaired. The node reaches this state if the system has a hardware or software failure.
INIT	In the initial state, the node is being initialized. A node stays in this state when it is coming up and remains in this state until it has been fully initialized. Being fully initialized means that all of the hardware has been initialized correctly and there are no diagnostic faults.
MASTER	In the primary state, the node is responsible for all switch management functions.
STANDBY	In the standby state, leader election occurs—the primary and backup nodes are elected. The priority of the node is only significant in the standby state.

Example

The following command displays the status of the node, the priority of the node, and the general health of the system:

```
show node
```

The following is sample output from this command:

```

Node    State      Priority    SwHealth    HwHealth
-----
1  MASTER          0         49         7
2  BACKUP         0 49         7
```

If you specify the detail option, the same information is displayed on a per node basis rather than in a tabular format, as shown in the following example:

```

Node 1 information:
Node State:    MASTER
Node Priority:  0
Sw Health:     49
Hw Health:     7
Node MSM-B information:
Node State:    BACKUP
Node Priority:  0
```

```
Sw Health:    49
Hw Health:    7
```

History

This command was first available in an ExtremeXOS 10.1.

Platform Availability

This command is available only on SummitStack.

show nodealias

```
show nodealias
```

Description

This command shows basic information collected by the Node Alias feature. Node Alias discovers information about the end systems on a per-port basis. Information from packets from end systems, such as VLANID, source MAC address, source IP address, protocol, etc. are captured in a database that can be queried.

Syntax Description

This command has no arguments or variables.

Default

None.

Usage Guidelines

The following information about the Node Alias feature appears:

- **Max Entries**—Total number of the entries learned through the Node Alias feature on the switch.
- **Active Entries**—Number of the active entries learned through the Node Alias feature on the switch.
- **Purge Time**—Last time an entry was purged from the Node Alias table on the switch.
- **State**—State of the Node Alias table on the switch.
- **Protocols Enabled**—Lists the currently enabled protocols that Node Alias detects.
- **Ports Enabled**—Lists the currently Node Alias-enabled ports with each port's active and maximum (configured using `configure nodealias ports [port_list |all] maxentries entries`) number of alias entries. This information appears in the form: port # (active alias entries #/maximum alias entries #). For example, "1(2/10)".

Example

The following example shows basic information about the Node Alias feature:

```
# show nodealias
Max Entries:          0                      Active Entries: 0
Purge Time:           State:                2
Protocols Enabled:    ip, ipv6, ospf, bgp, bootps, bootpc, vrrp, dhcp, dhcpc, bpdu,
udp, mdns, llmnr, ssdp,
Ports Enabled:
Shown in parentheses are the active # and maximum # aliases on the port.
```

History

This command was first available in ExtremeXOS 22.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show nodealias ip address

```
show nodealias ip ip_address
```

Description

This command shows alias information collected by the Node Alias feature for the specified IP address. Node Alias discovers information about the end systems on a per-port basis. Information from packets from end systems, such as VLANID, source MAC address, source IP address, protocol, etc. are captured in a database that can be queried.

Syntax Description

nodealias	Node Alias feature that maps source IP address, MAC address, host name, and protocol on a per port basis.
ip	Designates displaying node aliases with the specified IP address.
<i>ip_address</i>	Specifies for which IP address to display node alias information.

Default

N/A

Example

The following example shows node alias information for IP address 10.51.2.1:

```
# show nodealias ip 10.51.2.1
```

Port	MAC Address	Alias ID	Time Last Learned	VID	Protocol	Source IP
10	00:00:5e:00:01:02	716168949	2016-06-24 13:26:41	1	ip	10.51.2.1

History

This command was first available in ExtremeXOS 22.2.

Platform Availability

VLAN

show nodealias mac

```
show nodealias mac mac_address {protocol protocol_name | ports
    [port_list | all]}
```

Description

This command shows alias information collected by the Node Alias feature for the specified MAC address. Node Alias discovers information about the end systems on a per-port basis. Information from packets from end systems, such as VLANID, source MAC address, source IP address, protocol, etc. are captured in a database that can be queried.

Syntax Description

nodealias	Node Alias feature that maps source IP address, MAC address, host name, and protocol on a per port basis.
mac	Designates displaying node aliases with the specified MAC address.
<i>mac_address</i>	Specifies for which MAC address to display node alias information.
protocol	Designates that you want to see only alias information discovered by a particular protocol (default is all protocols).
<i>protocol_name</i>	Specifies the protocol that you want to see alias information discovered by: ip, ipv6, ospf, bgp, bootps, bootpc, vrrp, dhcps, dhcpc, bpdu, udp, mdns, llmnr, or ssdp.
ports	Designates that you want to only see node alias information for the selected ports (default is all ports).
<i>port_list</i>	Specifies for which ports to show node alias information. Designated as a port list separated by comma (,) or dash (-).
all	(Default) Shows node alias information for all ports.

Default

If a specific protocol is not selected, information appears for all protocols. If specific port(s) are not selected, information appears for all ports.

Example

The following example shows node alias information for MAC address 20:b3:99:5e:b5:e5 and on all ports, discovered by all protocols:

```
# show nodealias mac 20:b3:99:5e:b5:e5
```

Port	MAC Address	Alias ID	Time Last Learned	VID	Protocol	Source IP
10	20:b3:99:5e:b5:e5	716168878	2016-06-24 13:23:57	1	ip	10.50.0.2

History

This command was first available in ExtremeXOS 22.2.

Platform Availability

VLAN

show nodealias ports

```
show nodealias ports {port_list| all}
```

Description

This commands shows, for the specified ports, the aliases discovered using the Node Alias feature (MAC address, alias ID, time last learned, VID, protocol, and source IP address). Node Alias discovers information about the end systems on a per-port basis. Information from packets from end systems are captured in a database that can be queried.

Syntax Description

nodealias	Node Alias feature that maps source IP address, MAC address, host name, and protocol on a per port basis.
ports	Specifies information on the ports for the Node Alias feature.
<i>port_list</i>	Specifies for which ports to display information. Designated as a port list separated by comma (,) or dash (-).
all	(Default) Specifies showing Node Alias information for all ports.

Default

If the ports are not specified, information appears for all ports.

Example

The following example shows Node Alias information for port 10:

```
# show nodealias ports 10
```

Port	MAC Address	Alias ID	Time Last Learned	VID	Protocol	Source IP
----	-----	-----	-----	----	-----	

```

-----
10      5c:26:0a:57:e9:2a      716168369  2016-06-24 13:23:56  1      dhcpc
10      00:04:96:97:d1:86      716168370  2016-06-24 13:23:57  1      ip          10.50.121.30
10      20:b3:99:5e:b5:e5      716168371  2016-06-24 13:23:57  1      ip          10.50.4.2
10      20:b3:99:5e:b5:f5      716168378  2016-06-24 13:23:57  1      ipv6

```

History

This command was first available in ExtremeXOS 22.2.

Platform Availability

VLAN

show nodealias protocol

```
show nodealias protocol protocol_name
```

Description

This command shows alias information collected by the Node Alias feature by the specified protocol. Node Alias discovers information about the end systems on a per-port basis. Information from packets from end systems, such as VLANID, source MAC address, source IP address, protocol, etc. are captured in a database that can be queried.

Syntax Description

nodealias	Node Alias feature that maps source IP address, MAC address, host name, and protocol on a per port basis.
protocol	Designates that you want to see only alias information discovered by a particular protocol (default is all protocols).
<i>protocol_name</i>	Specifies the protocol that you want to see alias information discovered by: ip, ipv6, ospf, bgp, bootps, bootpc, vrrp, dhcps, dhcpc, bpdu, udp, mdns, llmnr, or ssdp.

Default

If a specific protocol is not selected, information appears for all protocols. If specific port(s) are not selected, information appears for all ports.

Example

The following example...

```
example in codeblock
```

History

This command was first available in ExtremeXOS 22.2.

Platform Availability

VLAN

show ntp

```
show ntp
```

Description

Displays the global NTP status of the switch.

Syntax Description

N/A.

Default

N/A.

Usage Guidelines

N/A.

Example

The following command shows the NTP status of the switch:

```
Switch# show ntp
NTP                : Enabled
Authentication     : Disabled
Broadcast-Client   : Disabled
VR                 : VR-Default
```

History

This command was first available in ExtremeXOS 12.7.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show ntp association statistics

```
show ntp association [{ip_address} | {host_name}] statistics
  [{vr}vr_name]
```

Description

Shows NTP-related statistics about a specific NTP server.

Syntax Description

<i>ip_address</i>	Server or peer IP address.
<i>host_name</i>	Server or peer name.
vr	Specifies information for a particular virtual router.
<i>vr_name</i>	Name of the particular virtual router. If not specified, VR of current command context is used.

Default

N/A.

Usage Guidelines

N/A.

Example

The following command shows NTP-related statistics about the NTP server called “1.us.pool.ntp.org”:

```
Switch# show ntp association 1.us.pool.ntp.org statistics
Remote Host       : 1.us.pool.ntp.org
Local Interface   : 10.45.203.74
Time Last Received : 40 second
Time Until Next Send: 27 second
Reachability Change : 849 second
Packets Sent      : 18
Packets Received   : 18
Bad Authentication : 0
Bogus Origin       : 0
Duplicate          : 0
Bad Dispersion     : 0
Bad Reference Time : 0
Candidate Order    : 4
Peer Flags         : Config, Broadcast Client, Initial Burst
```

The following command shows NTP-related statistics for the NTP server at 128.138.141.172 for virtual router “VR-Mgmt”:

```
# show ntp association 128.138.141.172 statistics vr "VR-Mgmt"
VR Name           : VR-Mgmt
Remote Host       : 128.138.141.172
Local Interface   : 10.127.2.180
Time Last Received : 146 seconds
Time Until Next Send: 113 seconds
Reachability Change : 2564 seconds
Packets Sent      : 86
Packets Received   : 86
Bad Authentication : 0
Bogus Origin       : 0
```

```
Duplicate      : 0
Bad Dispersion : 0
Bad Reference Time : 0
Candidate Order : 6
```

History

This command was first available in ExtremeXOS 12.7.

Virtual router keyword added in ExtremeXOS 22.2

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show ntp association

```
show ntp association [{ip_address} | {host_name}] [{vr} vr_name]
```

Description

Shows all of the NTP clock source information, from a statically configured server, peer, or broadcast server. The NTP service updates the local clock from only one NTP server, with the best stability and stratum value which is considered as a system peer.

Syntax Description

<i>ip_address</i>	Server or peer IP address.
<i>host_name</i>	Server or peer name.
vr	Specifies information for a particular virtual router.
<i>vr_name</i>	Name of the particular virtual router. If not specified, VR of current command context is used.

Default

N/A.

Usage Guidelines

N/A.

Example

The following command shows detailed information about the NTP server:

```
# show ntp association
VR Name      Remote                               Reference ID   St Poll Reach Delay   Offset   Disp
=====
=====
```

VR-Mgmt	!45.125.255.53	223.255.185.2	2	64	1	0.01172	-0.08789
1.98431							
VR-Mgmt	!10.127.2.255	BCST	16	64	0	0.00000	+0.00000
4.00000							

The following command shows detailed information for the NTP server at 128.138.141.172 for virtual router "VR-Mgmt":

```
# show ntp association 128.138.141.172 vr "VR-Mgmt"
VR Name           : VR-Mgmt
Remote IP         : 128.138.141.172 Local IP       : 10.127.2.180
Host Mode        : Client Peer Mode             : Server
Version          : 3 Key ID                    : 0
Stratum          : 1 Precision                 : -29
Leap             : 00 Reference ID              : .NIST.
Root Distance    : 0.00024 Root Dispersion     : 0.00049
Reachability     : 377 UnReachability          : 0
Peer Poll        : 13 Host Poll                : 8
Broadcast Offset : 0.23016 TTL/Mode             : 0
Offset           : -0.007505 Delay              : 0.23016
Error Bound      : 0.09737 Filter Error         : 0.04152
Peer Flags       : System Peer, Config
Reference Time   : db9e2e30.00000000 Tue, Oct 4 2016 13:33:36.000
Originate Time   : 00000000.00000000 Thu, Feb 7 2036 6:28:16.000
Receive TimeStamp : db9e2e32.e70200ee Tue, Oct 4 2016 13:33:38.902
Transmit TimeStamp : db9e2e32.e70200ee Tue, Oct 4 2016 13:33:38.902
Filter Order     : 2 3 7 0 6 4
1 5
Filter Delay     : 0.23088 0.23265 0.23016 0.23039 0.23238 0.23882 0.23152
0.23042
Filter Offset    : -0.00573 -0.00759 -0.00750 -0.00671 -0.00780 -0.00995 -0.00594
-0.00458
```

History

This command was first available in ExtremeXOS 12.7.

Virtual router keyword added in ExtremeXOS 22.2

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show ntp key

```
show ntp key
```

Description

Shows the NTP key index number, trusted or non-trusted, authentication type, and encrypted key string.

Syntax Description

N/A.

Default

N/A.

Usage Guidelines

N/A.

Example

The following command shows the NTP key index number, trusted or non-trusted, authentication type, and encrypted key string:

```
# show ntp key
Key Index      Trusted    Auth      Key String (encrypted)
=====
100             No         MD5       67:74:7d:78:6f:6c:67:5b:33
200             Yes        SHA-256   23:24:6c:35:4a:35:79:74:65
```

History

This command was first available in ExtremeXOS 12.7.

SHA-256 information was added in ExtremeXOS 30.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show ntp restrict-list

```
show ntp restrict-list {user | system} [{vr} vr_name]
```

Description

Shows the NTP access list of the current system based on the source IP blocks.

Syntax Description

user	Displays the NTP access list of the current user.
system	Displays the for the current system.
all	Displays both user and system data.
vr	Specifies showing NTP information for a given VR.
<i>vr_name</i>	Specifies which VR to show information about. If a VR name is not specified, the VR of the current command context is used.

Default

Displays all by default.

Usage Guidelines

N/A.

Example

The following command displays all NTP access list information for VR *"VR-Mgmt"*:

```
# show ntp restrict-list user vr "VR-Mgmt"
```

VR Name	IP Address	Mask	Count	Type	Action
VR-Mgmt	1.1.1.1	255.255.255.255	0	User	Permit

History

This command was first available in ExtremeXOS 12.7.

The **vr** keyword was added in ExtremeXOS 22.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show ntp server

```
show ntp server
```

Description

Shows the NTP servers configured on the switch, including the name, IP address, key ID, and index.

Syntax Description

N/A.

Default

N/A.

Usage Guidelines

N/A.

Example

The following command shows the NTP servers configured on the switch:

```
# show ntp server
```

VR Name	Name	IP Address	Type	Flags	Key
Index					

=====					
VR-Mgmt	45.125.255.53	45.125.255.53	Server	I	-

History

This command was first available in ExtremeXOS 12.7.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show ntp sys-info

```
show ntp sys-info
```

Description

Shows the current system status based on the most reliable clock server or NTP server.

Syntax Description

N/A.

Default

N/A.

Usage Guidelines

N/A.

Example

The following command shows the current system status based on the most reliable clock server or NTP server:

```
switch # show ntp sys-info
System Peer       : 0.us.pool.ntp.org
System Peer Mode  : Client
Leap Indicator    : 00
Stratum           : 3
Precision         : -20
Root Distance     : 0.09084 second
Root Dispersion   : 0.23717 second
Reference ID      : [216.93.242.12]
Reference time    : d140571d.e8389ff7  Fri, Apr  1 2011  6:52:29.907
System Flags      : Monitor, Ntp, Kernel, Stats
Jitter           : 0.004700 second
Stability         : 0.000 ppm
Broadcast Delay   : 0.007996 second
Auth Delay        : 0.000000 second
```

History

This command was first available in ExtremeXOS 12.7.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show ntp vlan

show ntp vlan {{**vr**} *vr_name*}

Description

Shows the NTP status of each VLAN configured on the switch.

Syntax Description

vr	Specifies showing NTP information for a given VR.
<i>vr_name</i>	Specifies which VR to show information about. If a VR name is not specified, the VR of the current command context is used.

Default

If a VR name is not specified, the VR of current command context is used.

Usage Guidelines

N/A.

Example

The following command shows the NTP status of VLAN "mgmt":

# show ntp vlan "mgmt"				
VR Name	Vlan	NTP Status	Broadcast Server	Key Index
=====				
VR-Mgmt	Mgmt	Disabled	Disabled	-

The following example displays all VLANs for VR "VR-Default":

# show ntp vlan vr "VR-Default"				
VR Name	Vlan	NTP Status	Broadcast Server	Key Index
=====				
VR-Default	Default	Disabled	Disabled	-
VR-Default	vlan1	Disabled	Disabled	-

History

This command was first available in ExtremeXOS 12.7.

The **vr** keyword was added in ExtremeXOS 22.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show ntp vr

```
show ntp {{vr}}vr_name}
```

Description

This command shows NTP information for the specified VR.

Syntax Description

ntp	Network Time Protocol.
vr	Specifies showing NTP information for a given VR.
<i>vr_name</i>	Specifies which VR to show information about. If a VR name is not specified, the VR of the current command context is used.

Default

If a VR name is not specified, the VR of the current command context is used.

Example

The following example shows NTP information for VR "vr1".

```
# show ntp
NTP           : Enabled
Authentication : Disabled
Broadcast-Client : Disabled
VR            : vr1
```

History

This command was first available in ExtremeXOS 22.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show odometers

```
show odometers
```

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 individual components in the switch have been functioning since it was manufactured. This odometer counter is kept in the EEPROM of each monitored component.

Recorded Statistics

The following odometer statistics are collected by the switch:

- Service Days—The amount of days that the component has been running.
- First Recorded Start Date—The date that the component was powered-up and began running.

Depending on the software version running on your switch and the type of switch you have, additional or different odometer information may appear.

Example

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

```
show odometers
```

The following is sample output from a stand-alone switch:

Service Field Replaceable Units	Days	Start Date	Service Field Replaceable Units	Days	Start Date
Switch : SummitX	7	Dec-08-2004	Switch : X(SS)	381	Oct-29-2009
			VIM1-SS-1 :	376	Jul-30-2009

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show ospf

```
show ospf
```

Description

Displays global *OSPF* information.

Syntax Description

This command has no keywords or arguments.

Default

N/A.

Usage Guidelines

Not applicable.

Example

The following command displays global OSPF information:

```
show ospf
```

The following is sample output from this command:

```

OSPF                : Disabled          MPLS LSP as Next-Hop: No
RouterId            : 0.0.0.0           RouterId Selection  : Automatic
InstanceId          : 22
ASBR                : No                ABR                : No
ExtLSA              : 0                 ExtLSAChecksum      : 0x0
OriginateNewLSA     : 0                 ReceivedNewLSA      : 0
SpfHoldTime         : 3                 Lsa Batch Interval  : 30s
CapabilityOpaqueLSA : Enabled
10M Cost            : 22                 100M Cost           : 20
1000M Cost (1G)     : 19                 2500M Cost (2.5G)   : 18
5000M Cost (5G)     : 17                 10000M Cost (10G)   : 16
25000M Cost (25G)   : 15                 40000M Cost (40G)   : 13
50000M Cost (50G)   : 12                 100000M Cost (100G) : 10
ASExternal LSALimit : Disabled           Timeout (Count)      : Disabled (0)
Originate Default   : Enabled             Always : Yes Type: 2  Cost: 10  Tag: 0
SNMP Traps          : Enabled             SNMP Trap Bit Map   : 0xffff
VXLAN Extensions    : Enabled
Redistribute:
Protocol            Status  cost  Type Tag      Policy
direct              Disabled 0     0    0          None
static              Disabled 0     0    0          None
rip                 Disabled 0     0    0          None

```

e-bgp	Disabled	0	0	0	None
i-bgp	Disabled	0	0	0	None
isis-level-1	Disabled	0	0	0	None
isis-level-2	Disabled	0	0	0	None
isis-level-1-external	Disabled	0	0	0	None
isis-level-2-external	Disabled	0	0	0	None
host-mobility	Enabled	0	2	0	None
# of Inter-VR Exports: 3					

History

This command was first available in ExtremeXOS 10.1.

The [SNMP Traps](#) and 40G parameters were added in ExtremeXOS 12.6.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

show ospf area

show ospf area {**detail** | *area-identifier*}

Description

Displays information about the [OSPF](#) area.

Syntax Description

detail	Specifies to display the information in detailed format.
<i>area-identifier</i>	Specifies an OSPF area.

Default

N/A.

Usage Guidelines

None.

Example

The following is sample output from this command:

# show ospf area									
AREA ID	Type	Summ	Def	Num	Num	SPF	Num	LSA	
Metric ABR	ASBR	Runs	LSAs	Checksum					
0.0.0.0	NORM	----	-----	0	0	0	0	0x0	

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [***Switch Engine 32.3 Feature License Requirements***](#) document.

show ospf ase-summary

```
show ospf ase-summary
```

Description

Displays the [*OSPF*](#) external route aggregation configuration.

Syntax Description

this command has no keywords or arguments.

Default

N/A.

Usage Guidelines

Not applicable.

Example

The following command displays the OSPF external route aggregation configuration:

```
show ospf ase-summary
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [***Switch Engine 32.3 Feature License Requirements***](#) document.

NEW! show ospf inter-vr-export

```
show ospf inter-vr-export
```

Description

Displays global information about *OSPF* route exports between VRs.

Syntax Description

This command has no keywords or arguments.

Default

N/A.

Usage Guidelines

Not applicable.

Example

The following command displays information about OSPF route exports:

```
# show ospf inter-vr-export

Route Type   Source VR      Cost  Type Tag      Flags Policy
ospf-intra   vr1000         -    -   -          -   vr1000pol
ospf-inter   vr1000         -    -   -          -   vr1000pol
ospf-externl virtrouter10> 60    1   10         X   None

> Indicates Source VR and Policy truncated past 12 characters.
Flags : (X) Exclude routes corresponding to private IP ranges.
```

History

This command was first available in ExtremeXOS 32.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

NEW!

show ospf inter-vr-export detail

```
show ospf inter-vr-export {detail}
```

Description

Displays detailed information *OSPF* route exports between VRs.

Syntax Description

detail	Specifies to display the information in detailed format.
---------------	--

Default

N/A.

Usage Guidelines

Not applicable.

Example

The following command displays detailed OSPF route exports information:

```
# show ospf inter-vr-export detail
Route Type      : ospf-intra
Source VR       : vr1000
Cost            : -
Type            : -
Tag             : -
Exclude Private : No
Policy          : vr1000pol

Route Type      : ospf-inter
Source VR       : vr1000
Cost            : -
Type            : -
Tag             : -
Exclude Private : No
Policy          : vr1000pol

Route Type      : ospf-extern1
Source VR       : virtrouter1001
Cost            : 60
Type            : 1
Tag             : 10
Exclude Private : Yes
Policy          : None
```

History

This command was first available in ExtremeXOS 32.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show ospf interfaces

```
show ospf interfaces {vlan vlan-name | area area-identifier | enabled}
```

Description

Displays information about one or all *OSPF* interfaces.

Syntax Description

<i>vlan-name</i>	Specifies a <i>VLAN</i> name.
<i>area-identifier</i>	Specifies an OSPF area.
enabled	Displays only OSPF enabled interfaces.

Default

If no argument is specified, all OSPF interfaces are displayed.

Usage Guidelines

Not applicable.

Example

The following command displays information about one or all OSPF interfaces on the VLAN accounting:

```
show ospf interfaces vlan accounting
```

The following output displays BFD protection configuration information:

```
# show ospf interfaces

VLAN      IP Address      AREA ID      Flags      Cost State  Neighbors
HQ_10_0_2  10.0.2.2        /24 0.0.0.0  -rifb--    4/A R      1
HQ_10_0_5  10.0.5.2        /24 0.0.0.0  -rif---    4/A BR     1

Flags:  b - BFD protection configured, f - Interface Forwarding Enabled,
        i - Interface OSPF Enabled, n - Multinetted VLAN,
        p - Passive Interface, r - Router OSPF Enable,
Cost:    A - Automatic Cost, C - Configured Cost.

Total number of interfaces: 2
```

The following output displays the BFD session state:

```
Interface(rif1000027): 10.0.2.2/24 Vlan: HQ_10_0_2 OSPF: ENABLED Router: ENABLED
AreaId: 0.0.0.0 RtId: 10.0.2.2 Link Type: broadcast(auto) Passive: No
Cost: 4/A Priority: 10 Transit Delay: 1 DAD State:Valid
Hello Interval: 10s Rtr Dead Time: 40s Retransmit Interval: 5s
Wait Timer: 40s
Authentication: NONE
State: DR Number of events: 1
DR RtId: 10.0.2.2 DR IP addr: 10.0.2.2 BDR IP addr: 10.0.2.1
Num Neighbor State Change to FULL : 1
BFD Protection: On

Neighbors:
  RtId: 10.0.3.1 IpAddr: 10.0.2.1 Pri: 5 Type: Auto
  State: FULL Dr: 10.0.2.2 BDR: 10.0.2.1 Dead Time: 00:00:00:03
  Options (0x42): Opaque LSA: Yes
  BFD Session State: Active
```

History

This command was first available in ExtremeXOS 10.1.

The enabled option was added in ExtremeXOS 12.2.

BFD display output was added in 15.3.2.

The D (duplicate address detected on VLAN) and T (tentative address) flags were removed in ExtremeXOS 30.1.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [***Switch Engine 32.3 Feature License Requirements***](#) document.

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

Not applicable.

Example

The following command displays information about all OSPF interfaces:

```
show ospf interfaces detail
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

show ospf lsdb

```
show ospf lsdb {detail | stats} {area [area-identifier | all]} [{lstype}
  [lstype | all]} {lsid lsid-address{lsid-mask}} {routerid routerid-
  address {routerid-mask}} {interface[[ip-address{ip-mask} | ipNetmask]
  | vlan vlan-name]}
```

Description

Displays a table of the current Link-State Database (LSDB).

Syntax Description

detail	Specifies to display all fields of matching LSAs in a multi-line format.
stats	Specifies to display the number of matching LSAs, but not any of their contents.
<i>area-identifier</i>	Specifies an OSPF area.
all	Specifies all OSPF areas.
<i>lstype</i>	Specifies an LS type.
lsid	Specifies an LS ID.
<i>lsid-mask</i>	Specifies an LS ID mask.
<i>routerid-address</i>	Specifies a LSA router ID address.
interface	Specifies to display interface types.
<i>vlan-name</i>	Specifies a VLAN name.

Default

Display in summary format.

Usage Guidelines

ExtremeXOS provides several filtering criteria for the show ospf lsdb command. You can specify multiple search criteria and only the 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 first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [***Switch Engine 32.3 Feature License Requirements***](#) document.

show ospf memory

```
show ospf memory {detail | memoryType}
```

Description

Displays *OSPF* specific memory usage.

Syntax Description

detail	Displays detail information.
<i>memoryType</i>	Specifies the memory type usage to display.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays OSPF specific memory for all types:

```
show ospf memory detail
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [***Switch Engine 32.3 Feature License Requirements***](#) document.

show ospf neighbor

```
show ospf neighbor {routerid [ip-address {ip-mask} | ipNetmask]} {vlan
    vlan-name} {detail}
```

Description

Displays information about an [*OSPF*](#) neighbor.

Syntax Description

<i>ip-address</i>	Specifies an IP address.
<i>ip-mask</i>	Specifies a subnet mask.
<i>ipNetmask</i>	Specifies IP address / Netmask.
<i>vlan-name</i>	Specifies a <i>VLAN</i> name.
detail	Specifies detail information.

Default

If no argument is specified, all OSPF neighbors are displayed.

Usage Guidelines

Not applicable.

Example

The following command displays information about the OSPF neighbors on the VLAN accounting:

```
show ospf neighbor vlan accounting
```

The following command output displays BFD protection status of all OSPF neighbors:

```
# show ospf neighbor
Neighbor ID      Pri State                Up/Dead Time      Address           Interface
      BFD Session State
=====
```

```

160.26.26.2      10 FULL      /BDR      10:16:42:57/00:00:00:00  160.26.26.2
CHI_160_26_26
    Disabled
10.0.2.2         10 FULL      /BDR      07:17:55:29/00:00:00:09  10.0.2.2      HQ_10_0_2
    Active
10.0.3.2         10 FULL      /BDR      07:17:54:56/00:00:00:03  10.0.3.2      HQ_10_0_3
    Error (Session Limit Exceeded)

Total number of neighbors: 3 (All neighbors in Full state)

# show ospf neighbor {vlan} <vlan-name>

Neighbor ID      Pri State                Up/Dead Time
Address          Interface
      BFD Session State
=====
10.0.3.2         1 FULL      /BDR      00:11:13:06/00:00:00:04  12.0.2.2      v2
    Active
Total number of neighbors: 1 (All neighbors in Full state)

# show ospf neighbor detail

Neighbor 10.0.3.2, interface address 12.0.2.2
  In the area 0.0.0.0 via interface v2
  Neighbor priority is 1,  State is INIT,38 state changes
  DR is 12.0.2.1 BDR is 12.0.2.2
  Options is 0x42
  Neighbor is up for 00:11:04:05
  Time since last Hello 00:00:00:00
  Retransmission queue length is 0
  BFD Session State: None

```

History

This command was first available in ExtremeXOS 10.1.

BFD output was added in ExtremeXOS 15.3.2.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [***Switch Engine 32.3 Feature License Requirements***](#) document.

show ospf virtual-link

```
show ospf virtual-link {router-identifier area-identifier}
```

Description

Displays virtual link information about a particular router or all routers.

Syntax Description

<i>router-identifier</i>	Specifies a router interface number.
<i>area-identifier</i>	Specifies an <u>OSPF</u> area.

Default

N/A.

Usage Guidelines

The area-identifier refer to the transit area used for connecting the two end-points. The transit area cannot have an area identifier of 0.0.0.0 and cannot be a stub or NSSA area.

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 first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the **Switch Engine 32.3 Feature License Requirements** document.

show ospfv3

```
show ospfv3
```

Description

Displays global OSPFv3 information.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays global OSPFv3 information:

```
show ospfv3
```

The following is sample output:

```

OSPFv3           : Disabled           RouterId         : 0.0.0.0
RouterId Selection : Automatic         ASBR            : No
ABR              : No                 ExtLSAs         : 0
ExtLSAChecksum   : 0x0               OriginateNewLSAs : 0
ReceivedNewLSAs  : 0                 SpfHoldTime     : 3s
Num of Areas     : 1                 LSA Batch Interval : 0s
10M Cost         : 100               100M Cost       : 50
1000M Cost (1G)  : 40                2500M Cost (2.5G) : 40
5000M Cost (5G)  : 40                10000M Cost (10G) : 20
25000M Cost (25G) : 20               40000M Cost (40G) : 20
50000M Cost (50G) : 20              100000M Cost (100G) : 10
Graceful Restart  : None              Grace Period     : 120s
Import Policy File : none
SNMP Traps       : Disabled

Redistribute:
  Protocol      Status   Cost   Type  Tag   Policy
  direct        Disabled 20      2    ---   none
  e-bgp          Disabled 20      2    ---   none
  i-bgp          Disabled 20      2    ---   none
  ripng          Disabled 20      2    ---   none
  static         Disabled 20      2    ---   none
  isis-level-1   Disabled 20      2    ---   none
  isis-level-2   Disabled 20      2    ---   none
  isis-level-1-external Disabled 20      2    ---   none
  isis-level-2-external Disabled 20      2    ---   none
  host-mobility  Disabled 20      2    ---   none

```

History

This command was first available in ExtremeXOS 11.2.

The 40G parameter was added in ExtremeXOS 12.6.

SNMP trap status information was added in ExtremeXOS 22.1.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [***Switch Engine 32.3 Feature License Requirements***](#) document.

show ospfv3 area

```
show ospfv3 area {area_identifier | detail}
```

Description

Displays information about [*OSPFv3*](#) areas.

Syntax Description

<i>area_identifier</i>	Specifies an OSPFv3 area, a four-byte, dotted decimal number.
detail	Specifies to display the information in detailed format.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays summary information about the OSPFv3 areas:

```
show ospfv3 area
```

The following is sample output:

AREA ID	Type	Summ	Def	Num	Num	Num	SPF	Num	LSA
Metric ABR	ASBR	Intf	Runs	LSAs	Checksum				
0.0.0.0	NORM	----	-----	0	0	1	7	7	0x3155b
1.0.0.0	NORM	----	-----	1	1	1	6	9	0x4793d
2.0.0.0	NORM	----	-----	0	0	1	5	10	0x47174
3.0.0.0	NORM	----	-----	1	0	1	3	12	0x420cf
5.0.0.0	NORM	----	-----	1	0	1	4	10	0x3b5b1

The following command displays information about OSPFv3 area 1.0.0.0:

```
show ospfv3 area 1.0.0.0
```

The following is sample output:

```

Area Identifier       : 1.0.0.0                Type                : NORM
Router ID            : 20.0.0.1                Num of Interfaces   : 1
Spf Runs              : 6                      Num ABRs            : 1
Num ASBRs             : 1                      Num DC-Bit LSAs     : 1
Num Indication LSAs   : 1                      Num of DoNotAge LSAs: 1
Num LSAs              : 9                      LSA Chksum          : 0x4793d
Num of Nbrs           : 1                      Num of Virtual Nbrs : 0
Interfaces:
Interface Name        Ospf State   DR ID              BDR ID
to65                  E    BDR             0.0.0.65           20.0.0.1
accounts              E    DR              80.0.0.5           0.0.0.0
finance              E    BDR             90.0.0.7           66.0.0.4
engineering           E    ODR             192.168.0.1        165.0.0.3
Corporate             E    ODR             201.0.16.6         204.0.0.1
Inter-Area route Filter: ospfSummPolicy
External route Filter: ospfExtPolicy
Configured Address Ranges:
Addr: fffe:408:1449::/48 Type: 3 Advt: Yes
Addr: ffe0:930:2781::/40 Type: 7 Advt: No

```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [*Switch Engine 32.3 Feature License Requirements*](#) document.

show ospfv3 interfaces

```
show ospfv3 interfaces {vlan vlan_name | tunnel tunnel_name | area
area_identifier | detail}
```

Description

Displays information about one or all [*OSPFv3*](#) interfaces.

Syntax Description

<i>vlan_name</i>	Specifies an IPv6 configured <i>VLAN</i> .
<i>tunnel_name</i>	Specifies an IPv6 tunnel.
<i>area_identifier</i>	Specifies an OSPFv3 area, a four-byte, dotted decimal number.
detail	Specifies to display the information in detailed format.

Default

If no argument is specified, all OSPFv3 interfaces are displayed.

Usage Guidelines

None.

Example

The following command shows a summary of the OSPFv3 interfaces:

```
show ospfv3 interfaces
```

The following is sample output from the command:

```
Interface  IPV6 Address          AREA ID      Flags Cost    State  Nbrs
v100       2001:db8:10::2/64      0.0.0.0      -rif- 40/A   P2P    1
v101       2001:db8:20::1/64      1.0.0.0      -rif- 40/A   BDR    1

Flags : (b) BFD protection configured, (f) Interface Forwarding Enabled,
        (i) Interface OSPFv3 Enabled, (p) Passive Interface, (r) Router OSPFv3 Enable.
Cost   : (A) Automatic cost, (C) Configured cost.
```

The following command displays information about the OSPFv3 interface v100:

```
show ospfv3 interfaces "v100"
```

The following is sample output:

```
Interface          : v100                Enabled          : ENABLED
Router             : ENABLED             AreaID          : 0.0.0.0
RouterID           : 10.1.1.2            Link Type       : point-to-point
Passive            : No                  Cost            : 40/A
Priority            : 1                   Transit Delay    : 1s
Hello Interval     : 10s                 Rtr Dead Time   : 40s
Retransmit Interval : 5s                 Wait Timer      : 40s
Interface ID       : 19                  Instance ID     : 0
State              : P2P                 Number of state chg : 1
Hello due in       : 7s                 Number of events  : 2
Total Num of Nbrs  : 1                   Nbrs in FULL State : 1
Hellos Rxed        : 127733             Hellos Txed      : 127739
DB Description Rxed : 4                  DB Description Txed : 3
LSA Request Rxed   : 1                   LSA Request Txed  : 1
LSA Update Rxed    : 2121                LSA Update Txed   : 6156
LSA Ack Rxed       : 5962                LSA Ack Txed      : 2121
In Discards        : 0
DR RtId            : 0.0.0.0             BDR RtId        : 0.0.0.0
BFD Protection     : Off
Authentication     : IPsec ESP
SPI                : 256                 Algorithm       : HMAC-SHA-256
Encrypted Key String: #$bC0GrQHhqeDABqYwnlf dihIsOhV9g==

Neighbors:
  RtrId: 10.1.1.1 IpAddr: fe80::204:96ff:fe51:ea8e Pri: 1 Type: Auto
  State: FULL DR: 0.0.0.0 BDR: 0.0.0.0 Dead Time: 00:00:31
  Options: 0x13 (-|R|-|-|E|V6) Opaque LSA: No
  BFD Session State: None
```

The following command displays information about the OSPFv3 interface intf1 using Authentication Trailer (line 24):

```
# show ospfv3 interfaces intf1

Interface          : intf1                Enabled          : ENABLED
Router             : ENABLED             AreaID          : 0.0.0.0
RouterID           : 0.0.0.2            Link Type       : broadcast
Passive            : No                  Cost            : 40/A
Priority            : 1                   Transit Delay    : 1s
Hello Interval     : 10s                 Rtr Dead Time   : 40s
Retransmit Interval : 5s                 Wait Timer      : 40s
Interface ID       : 64                  Instance ID     : 0
State              : DR                 Number of state chg : 2
Hello due in       : 9s                 Number of events  : 5
Total Num of Nbrs  : 1                   Nbrs in FULL State : 1
Hellos Rxed        : 50                 Hellos Txed      : 51
DB Description Rxed : 3                  DB Description Txed : 3
LSA Request Rxed   : 1                   LSA Request Txed  : 1
LSA Update Rxed    : 2                   LSA Update Txed   : 6
LSA Ack Rxed       : 4                   LSA Ack Txed      : 1
In Discards        : 0
DR RtId            : 0.0.0.2             BDR RtId        : 0.0.0.1
Restart Helper     : None
Restart Helper Strict LSA Checking: Enabled
BFD Protection     : Off
Authentication     : Keychain auth1
```

```

Neighbors:
  RtrId: 0.0.0.1  IpAddr: fe80::204:96ff:fe97:e9e1  Pri: 1  Type: Auto
  State: FULL  DR: 0.0.0.2  BDR: 0.0.0.1  Dead Time: 00:00:36
  Options: 0x13 (-|R|-|-|E|V6)
  BFD Session State: None

```

The following command output shows BFD protection configuration information:

```

# show ospfv3 interfaces
Interface  IPV6 Address          AREA ID      Flags Cost   State   Nbrs
HQ_10_0_4  2000::d00:202/64       0.0.0.0     -rifb 40/A   DR      1
Flags : (b) BFD protection configured, (f) Interface Forwarding Enabled,
        (i) Interface OSPFv3 Enabled, (p) Passive Interface, (r) Router OSPFv3 Enable.
Cost   : (A) Automatic cost, (C) Configured cost.

```

The following command output displays the BFD session state:

```

# show ospfv3 interfaces {vlan} <vlan-name>
Interface      : HQ_10_0_4      Enabled      : ENABLED
Router         : ENABLED       AreaID       : 0.0.0.0
RouterID       : 2.2.2.2       Link Type    : broadcast
Passive        : No           Cost         : 40/A
Priority        : 1            Transit Delay : 1s
Hello Interval : 10s          Rtr Dead Time : 40s
Retransmit Interval : 5s      Wait Timer    : 40s
Interface ID   : 50           Instance ID   : 0
State          : DR           Number of state chg : 2
Hello due in   : 4s          Number of events : 4
Total Num of Nbrs : 1        Nbrs in FULL State : 1
Hellos Rxed    : 1306       Hellos Txed    : 1306
DB Description Rxed : 5      DB Description Txed : 3
LSA Request Rxed : 1        LSA Request Txed : 1
LSA Update Rxed  : 18       LSA Update Txed  : 37
LSA Ack Rxed     : 36       LSA Ack Txed     : 17
In Discards     : 0
DR RtId         : 2.2.2.2     BDR RtId      : 1.1.1.1
DR Interface addr : fe80::201:30ff:fe10:3b16
BDR Interface addr : fe80::201:30ff:fe10:3ae6
BFD Protection   : On
Authentication   : IPsec ESP
SPI             : 256         Algorithm       : HMAC-SHA-256
Encrypted Key String: #$bC0GrQHhqeDABqYwnlfdihIsOhV9g==

Neighbors:
  RtrId: 1.1.1.1  IpAddr: fe80::201:30ff:fe10:3ae6  Pri: 1  Type: Auto
  State: FULL  DR: 2.2.2.2  BDR: 1.1.1.1  Dead Time: 00:00:40
  Options: 0x13 (-|R|-|-|E|V6)  Opaque LSA: No
  BFD Session State: Pending

```

History

This command was first available in ExtremeXOS 11.2.

BFD example output was added in 15.3.2.

Support for new **Link Type** and **State** values (Link Type: point-to-point, State: P2P) were added in ExtremeXOS 15.7.1.

IPsec Authentication information was added in ExtremeXOS 31.2.

Authentication Trailer information was added in ExtremeXOS 31.3.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

show ospfv3 lsdb stats

```
show ospfv3 lsdb stats {area [area_identifier | all] {lstype [router
| network | inter-prefix | inter-router | intra-prefix | nssa]}
| [vlan [vlan_name | all] | tunnel [tunnel_name | all]] {lstype
link} | lstype [as-external | router | network | inter-prefix |
inter-router | intra-prefix | link]} {lsid lsid_address} {adv-router
router_identifier}
```

Description

Displays a table of the current Link-State Database (LSDB) statistics.

Syntax Description

<i>area_identifier</i>	Specifies an OSPFv3 area, a four-byte, dotted decimal number.
all	Specifies all OSPFv3 areas, IPv6 configured VLANs, or IPv6 tunnels.
<i>vlan_name</i>	Specifies an IPv6 configured VLAN .
<i>tunnel_name</i>	Specifies an IPv6 tunnel.
link	Link LSA.
router	Router LSA.
network	Network LSA.
inter-prefix	Inter Area Prefix LSA.
inter-router	Inter Area Router LSA.
intra-prefix	Intra Area Prefix LSA.
nssa	NSSA LSA.
as-external	AS External LSA.
<i>lsid_address</i>	Specifies the link state ID of the LSA.
<i>router_identifier</i>	Specifies the router identifier of the advertising router.

Default

Display in summary format.

Usage Guidelines

ExtremeXOS provides several filtering criteria for the `show ospfv3 lsdb stats` command. You can specify multiple search criteria and only the 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 ospfv3 lsdb stats
```

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.

Example

The following command displays all areas and all types in a summary format:

```
show ospfv3 lsdb stats
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [*Switch Engine 32.3 Feature License Requirements*](#) document.

show ospfv3 lsdb

```
show ospfv3 lsdb {detail} {area [area_identifier | all] {lstype [router | network | inter-prefix | inter-router | intra-prefix | nssa]} | [vlan [vlan_name | all] | tunnel [tunnel_name | all]] {lstype link} | lstype [as-external | router | network | inter-prefix | inter-router | intra-prefix | link]} {lsid lsid_address} {adv-router router_identifier}
```

Description

Displays a table of the current Link-State Database (LSDB).

Syntax Description

detail	Specifies to display all fields of matching LSAs in a multi-line format.
<i>area_identifier</i>	Specifies an <i>OSPFv3</i> area, a four-byte, dotted decimal number.

<i>vlan_name</i>	Specifies an IPv6 configured VLAN .
<i>tunnel_name</i>	Specifies an IPv6 tunnel.
all	Specifies all OSPFv3 areas, IPv6 configured VLANs, or IPv6 tunnels.
link	Link LSA
as-external	AS External LSA
router	Router LSA
network	Network LSA
inter-prefix	Inter Area Prefix LSA
inter-router	Inter Area Router LSA
intra-prefix	Intra Area Prefix LSA
nssa	NSSA LSA.
<i>lsid_address</i>	Specifies the link state ID of the LSA.
<i>router_identifier</i>	Specifies the router identifier of the advertising router.

Default

Display in summary format.

Usage Guidelines

ExtremeXOS provides several filtering criteria for the `show ospfv3 lsdb` command. You can specify multiple search criteria and only the 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 ospfv3 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 ospfv3 lsdb
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

show ospfv3 neighbor

```
show ospfv3 neighbor {routerid ip_address} {vlan vlan_name | tunnel
    tunnel_name} {detail}
```

Description

Displays information about an [OSPFv3](#) neighbor.

Syntax Description

<i>ip_address</i>	Specifies a neighbor router ID.
<i>vlan_name</i>	Specifies an IPv6 configured VLAN .
<i>tunnel_name</i>	Specifies an IPv6 tunnel.
detail	Specifies detail information.

Default

If no argument is specified, all OSPFv3 neighbors are displayed.

Usage Guidelines

None.

Example

The following command displays information about the OSPFv3 neighbors on the VLAN accounting:

```
show ospfv3 neighbor vlan accounting
```

The following three examples provide sample outputs for `show ospfv3 neighbor`:

```
# show ospfv3 neighbor

Neighbor ID      Pri State           Up/Dead Time      Interface          InstanceID
-----
          BFD Session State
=====
1.1.1.1          1  FULL    /BDR    00:03:40:45/00:00:38 HQ_10_0_4         0
          Active

# show ospfv3 neighbor detail

Neighbor 1.1.1.1, Interface address fe80::201:30ff:fe10:3ae6
  In the area 0.0.0.0 via interface HQ_10_0_4
  Neighbor priority is 1,  State is FULL, 1338 events, 6 state changes
```

```
DR is 2.2.2.2 BDR is 1.1.1.1
Options is 0x13 (-|R|-|-|E|V6)
Neighbor is up for 00:03:42:17
Neighbor will be dead in 00:00:37
Retransmission queue length is 0
BFD Session State: Active

# show ospfv3 neighbor {vlan} <vlan-name>

Neighbor ID      Pri State           Up/Dead Time      Interface          InstanceID
-----
          BFD Session State
=====
1.1.1.1          1  FULL    /BDR             00:20:37:17/00:00:39 HQ_10_0_4         0
    Active
3.3.3.3          1  FULL    /DR              00:20:37:17/00:00:39 HQ_10_0_4         0
    Active
4.4.4.4          1  2WAY    /DOTHER          00:20:37:17/00:00:39 HQ_10_0_4         0
    None
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [*Switch Engine 32.3 Feature License Requirements*](#) document.

show ospfv3 virtual-link

```
show ospfv3 virtual-link {{routerid} router_identifier {area}
area_identifier}
```

Description

Displays virtual link(s) information.

Syntax Description

<i>router_identifier</i>	Specifies a router identifier, a four-byte, dotted decimal number.
<i>area_identifier</i>	Specifies an <i>OSPFv3</i> area, a four-byte, dotted decimal number.

Default

N/A.

Usage Guidelines

Router-identifier—Router ID for the other end of the link.

Area-identifier—Transit area used for connecting the two end-points. The transit area cannot have an area identifier of 0.0.0.0 and cannot be a stub or NSSA area.

Example

The following command displays information about the virtual link to a particular router:

```
# show ospfv3 virtual-link

Neighbor Router ID : 0.0.0.3          Transit Area ID      : 0.0.0.1
Hello Interval      : 10s             Router Dead Interval: 40s
Retransmit Interval: 5s              Transit Delay        : 1s
Instance ID        : 0               Via Interface        : intf2
State              : P2P              Num of State Changes: 1
Hello Due in       : 2               Number of Events     : 2
Hellos Rxed       : 149             Hellos Txed          : 148
DB Description Rxed : 2              DB Description Txed  : 2
LSA Request Rxed   : 1              LSA Request Txed    : 1
LSA Update Rxed    : 2              LSA Update Txed     : 2
LSA Ack Rxed       : 2              LSA Ack Txed        : 1
In Discards        : 0
Local Address       : 2001:20::1
Remote Address      : 2001:20::2
Restart Helper      : None
Restart Helper Strict LSA Checking: Enabled
Authentication      : IPsec ESP
SPI                 : 500             Algorithm            : HMAC-SHA-512
Encrypted Key String: #$bC0GrQHhqeDABqYwnlfidihIsOhV9g==

Virtual Neighbor:
  RtrId: 0.0.0.3  IpAddr: 2001:20::2  Type: Auto
  State: FULL  DR: 0.0.0.0  BDR: 0.0.0.0  Dead Time: 38
  Options: 0x13 (-|R|-|-|E|V6)
```

The following command displays information about the OSPFv3 virtual link authentication with default (none) configuration (line 19):

```
# show ospfv3 virtual-link

Neighbor Router ID : 0.0.0.3          Transit Area ID      : 0.0.0.1
Hello Interval      : 10s             Router Dead Interval: 40s
Retransmit Interval: 5s              Transit Delay        : 1s
Instance ID        : 0               Via Interface        : intf2
State              : P2P              Num of State Changes: 1
Hello Due in       : 2               Number of Events     : 2
Hellos Rxed       : 149             Hellos Txed          : 148
DB Description Rxed : 2              DB Description Txed  : 2
LSA Request Rxed   : 1              LSA Request Txed    : 1
LSA Update Rxed    : 2              LSA Update Txed     : 2
LSA Ack Rxed       : 2              LSA Ack Txed        : 1
In Discards        : 0
Local Address       : 2001:20::1
Remote Address      : 2001:20::2
Restart Helper      : None
Restart Helper Strict LSA Checking: Enabled
Authentication      : None

Virtual Neighbor:
  RtrId: 0.0.0.3  IpAddr: 2001:20::2  Type: Auto
  State: FULL  DR: 0.0.0.0  BDR: 0.0.0.0  Dead Time: 38
  Options: 0x13 (-|R|-|-|E|V6)
```

History

This command was first available in ExtremeXOS 11.2.

IPsec Authentication information was added in ExtremeXOS 31.2.

Authentication Trailer information was added in ExtremeXOS 31.3.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [*Switch Engine 32.3 Feature License Requirements*](#) document.

show pim anycast-rp

```
show pim {ipv4 | ipv6} anycast-rp {ip_address}
```

Description

Shows rendezvous point (RP) for Anycast RP using PIM (*RFC 4610*) information.

Syntax Description

ipv4	Specifies IPv4 address family.
ipv6	Specifies IPv6 address family.
anycast-rp	Shows Anycast RP information.
<i>ip_address</i>	Shows information for a specific Anycast RP.

Default

N/A.

Usage Guidelines

The Anycast RP using PIM feature provides fast convergence when RP routers fail using PIM protocol without using the source discovery protocol Multicast Source Discovery Protocol (MSDP) for both IPv4 and IPv6 address families.

To configure or remove RPs, use the `configure pim {ipv4 | ipv6} anycast-rp ip_address [policy | none]` command.

Example

The following example shows information for all Anycast RP:

```
# show pim anycast-rp
Anycast RP      Policy
-----
Peer's Address  Reg/Stop In  Reg/Stop Out
-----
```

```

150.150.150.1      pim-domain1-v4-anycast-rp
                   10.10.10.1*          50          50
                   20.20.20.1          45          100
                   30.30.30.1          45          56
100.100.100.1     pim-domain1-v4-anycast-rp
                   10.10.10.1*          0           0
                   20.20.20.1          0           0
                   30.30.30.1          0           0

Total number of anycast RP: 2
* Indicates local address

```

The following example shows information for all Anycast RP at IPv6 addresses:

```

# show pim ipv6 anycast-rp
Anycast RP
Peer's Address
-----
2002:db8:85a3::8a2e:370:7334      pim-domain1-v6-anycast-rp
2001:db8:85a3::8a2e:370:7334*      50          50
2001:db8:85a3::8a2e:370:7335      45          100
2001:db8:85a3::8a2e:370:7336      45          56
2001:db8:85a3::8a2e:370:7337      45          56
3002:db8:85a3::8a2e:370:7334      pim-domain2-v6-anycast-rp
3001:db8:85a3::8a2e:370:7334*      50          50
3001:db8:85a3::8a2e:370:7335      45          100
3001:db8:85a3::8a2e:370:7336      45          56

Total number of anycast RP: 2
* Indicates local address

```

The following example shows information for the Anycast RP at IPv6 address 2001:0db8:85a3:0000:0000:8a2e:0370:7334:

```

# show pim ipv6 anycast-rp 2002:0db8:85a3:0000:0000:8a2e:0370:7334
Anycast RP
Peer's Address
-----
2002:db8:85a3::8a2e:370:7334      pim-domain1-v6-anycast-rp
2001:db8:85a3::8a2e:370:7334*      50          50
2001:db8:85a3::8a2e:370:7335      45          100
2001:db8:85a3::8a2e:370:7336      45          56
2001:db8:85a3::8a2e:370:7337      45          56

Total number of anycast RP: 2
* Indicates local address

```

History

This command was first available in ExtremeXOS 31.1.

Platform Availability

This command is available on platforms that support the appropriate license for the PIM feature. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see the [Switch Engine 32.3 Feature License Requirements](#).

show pim cache

```
show pim { ipv4 | ipv6 } cache {{detail} | {state-refresh} {mlag-peer-info} {group_addr {source_addr}}}
```

Description

Displays the multicast cache entries created by PIM.

Syntax Description

ipv4	Displays IPv4 PIM information
ipv6	Displays IPv6 PIM information.
detail	Specifies to display the information in detailed format.
<i>group_addr</i>	Specifies an IP group address.
<i>source_addr</i>	Specifies an IP source address.
state-refresh	Specifies to display the PIM cache entries with state refresh parameters.
mlag-peer-info	Shows peer related information.

Default

N/A.

Usage Guidelines

Displays the following information:

- IP group address
- IP source address / source mask
- Upstream neighbor (RPF neighbor)
- Interface (VLAN) to upstream neighbor
- Cache expire time
- Egress and prune interface list

When the **detail** option is specified, the switch displays the egress VLAN list and the pruned VLAN list.

Example

The following command displays the PIM cache entry for group 239.255.255.1:

```
Switch.33 # show pim cache 239.255.255.1
Index  Dest Group      Source      InVlan      Origin
[0000] 239.255.255.1    124.124.124.124 (WR) v4      Sparse
Entry timer is not run; UpstNbr: 200.124.124.24
EgressIfList = vbs15(0) (FW) (SM) (I)
[0001] 239.255.255.1    118.5.1.1 (S)      vbs5      Sparse
```

```

Expires after 186 secs UpstNbr: 0.0.0.0
RP: 124.124.124.124 via 200.124.124.24 in v4
EgressIfList = vbs15(0) (FW) (SM) (I) , vpim5(170) (FW) (SM) (S)
PrunedIfList = v4(0) (SM)
Number of multicast cache = 20
Entry flags :-
R: RP tree. S: Source tree. W: Any source.
Egress/Pruned interface flags :-
SM: Sparse Mode          DM: Dense Mode
FW: Forwarding           PP: Prune pending
AL: Assert Loser         N: Neighbor present
I: IGMP member present   S: (s,g) join received
Z: (*,g) join received   Y: (*,*,rp) join received

```

The following command displays the PIM-DM cache entry with state-refresh information for group 225.0.0.1:

```

Switch.5 # show pim cache state-refresh 225.0.0.1
Index  Dest Group      Source              InVlan  Origin
[0001] 225.0.0.1      64.1.1.100 (S)      vixia   Dense   Not Pruned
Expires after 204 secs UpstNbr: 0.0.0.0
Refresh State: Originator(20), TTL: 16
EgressIfList = v36(0) (FW) (DM) (N)
[0001] 225.0.0.1      65.1.1.100 (S)      vixia   Dense   Not Pruned
Expires after 195 secs UpstNbr: 65.1.1.200
Refresh State: Not-Originator(25), TTL: 8
EgressIfList = v36(0) (FW) (DM) (N)

```

The following command displays the ingress VLAN information of all MLAG peers:

```

* (pacman debug) sw6.2 # show pim c mlag-peer-info
Index  Dest Group      Source              InVlan  Origin
[0000] 226.1.1.1      61.2.2.2 (WR)      fifteenth Sparse
Entry timer is not run; UpstNbr: 51.15.15.2
Peer Ingress VLAN (ISC 1): 51.15.15.4/24 (Same)
EgressIfList = eight(0) (FW) (SM) (I) , five(0) (FW) (SM) (I) , ten(0) (FW) (SM) (I)

[0001] 226.1.1.1      112.2.2.202 (S)     fifteenth Sparse
Expires after 203 secs UpstNbr: 51.15.15.2
RP: 61.2.2.2 via 51.15.15.2 in fifteenth
Peer Ingress VLAN (ISC 1): 51.15.15.4/24 (Same)
EgressIfList = eight(0) (FW) (SM) (I) , five(0) (FW) (SM) (I)
PrunedIfList = ten(0) (SM) (AL)

```

History

This command was first available in ExtremeXOS 10.1.

This command was modified in ExtremeXOS 15.2 to display MLAG peer information.

The **ipv4** and **ipv6** keyword options were added in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

show pim

```
show pim {ipv4 | ipv6 | detail | rp-set {group_addr} | vlan vlan_name}
```

Description

Displays the PIM configuration and statistics.

Syntax Description

ipv4	Displays PIM IPv4 configuration information.
ipv6	Displays PIM IPv6 configuration information.
detail	Displays show output in the detailed format.
<i>group_addr</i>	Specifies an IP multicast group, for which the RP is to be displayed.
<i>vlan_name</i>	Specifies a <u>VLAN</u> name.

Default

IPv4 is the default for the `show pim {ipv4 | ipv6}` command.

If no VLAN is specified, the configuration is displayed for all PIM interfaces.

If no multicast group is specified for the `rp-set` option (Rendezvous Point set), all RPs are displayed.

Usage Guidelines

The detail version of this command displays the global statistics for PIM, as well as the details of each PIM enabled VLAN.

Example

The following command displays the global PIM configuration and statistics:

```
* sw4.30 # show pim
PIM Enabled, Version 2
PIM CRP Disabled
BSR state           : ACCEPT_PREFERRED ; BSR Hash Mask : 255.255.255.252
Current BSR Info    : 61.2.2.2 (Priority 20) expires after 78 sec
Configured BSR Info : 0.0.0.0 (Priority 0)
CRP Adv Interval    : 60 sec ; CRP Holdtime: 150
BSR Interval        : 60 sec ; BSR Timeout : 130
Cache Timer         : 210 sec ; Prune Timer : 210
Assert Timeout      : 210 sec ; Register Suppression Timeout,Probe: 60, 5
Generation Id       : 0x52af433d      Dense Neighbor Check : On
PIM-DM State Refresh TTL           : 16
PIM-DM State Refresh Source Active Timer : 210
PIM-DM State Refresh Origination Interval : 60
Threshold for Last Hop Routers: 0 kbps
Threshold for RP                : 0 kbps
Register-Rate-Limit-Interval    : Always active
PIM SSM address range           : None
PIM Register Policy              : None
```



```

PIM Register Policy RP      : None
PIM IP Route Sharing        : Disabled
PIM IP Route Sharing Hash   : Source-Group-Next Hop
Register Checksum to exclude data
Active Sparse Ckts 10 Dense Ckts 0 State Refresh Ckts 0

Global Packet Statistics ( In          Out          Drop )
C-RP-Advs                   0             0             0
Registers                   0             0             0
RegisterStops               0             0             0

VLAN      Cid  IP Address      Designated      Flags      Hello J/P  Nbrs
              Router      Int    Int
eight      1  51.8.8.4      / 24 51.8.8.6    rifms----- 30    60    2
fifteenth  2  51.15.15.4     / 24 51.15.15.4  rifms----- 30    60    0
Legend: J/P Int: Join/Prune Interval
Flags : r - Router PIM Enabled, i - Interface PIM Enabled, f - Interface,
Forwarding Enabled, m - Interface Multicast Forwarding Enabled,
s - Sparse mode, d - Dense mode, c - CRP enabled,
t - Trusted Gateway configured, n - Multinetted VLAN,
p - Passive Mode, S - Source Specific Multicast, b - Border,
R - State Refresh Enabled.

```

The following command displays the detailed PIM configuration and statistics:

```

sw4.3 # show pim detail
PIM Enabled, Version 2
PIM CRP Disabled
BSR state      : ACCEPT_ANY ; BSR Hash Mask : 255.255.255.252
Current BSR Info : 0.0.0.0 (Priority 0)
Configured BSR Info : 0.0.0.0 (Priority 0)
CRP Adv Interval : 60 sec ; CRP Holdtime: 150
BSR Interval     : 60 sec ; BSR Timeout : 130
Cache Timer      : 210 sec ; Prune Timer : 210
Assert Timeout   : 210 sec ; Register Suppression Timeout,Probe: 60, 5
Generation Id    : 0x533331c7      Dense Neighbor Check : On
PIM-DM State Refresh TTL      : 16
PIM-DM State Refresh Source Active Timer : 210
PIM-DM State Refresh Origination Interval : 60
Threshold for Last Hop Routers: 0 kbps
Threshold for RP              : 0 kbps
Register-Rate-Limit-Interval : Always active
PIM SSM address range         : None
PIM Register Policy           : None
PIM Register Policy RP        : None
PIM IP Route Sharing          : Disabled
PIM IP Route Sharing Hash     : Source-Group-Next Hop
Register Checksum to exclude data
Active Sparse Ckts 10 Dense Ckts 0 State Refresh Ckts 0

Global Packet Statistics ( In          Out          Drop )
C-RP-Advs                   0             0             0
Registers                   0             0             0
RegisterStops               0             0             0

PIM SPARSE Interface[1] on VLAN eight is enabled and up
IP adr: 51.8.8.4      mask: 255.255.255.0      DR of the net: 51.8.8.6
DR Priority              : 1
Passive                  : No
Hello Interval           : 30 sec
Neighbor Time out       : 105 sec
Join/Prune Interval      : 60 sec
Join/Prune holdtime      : 210 sec
Trusted Gateway          : none

```

```

CRP group List      : none with priority 0
Shutdown priority   : 1024
Source Specific Multicast : Disabled
State Refresh       : Off
State Refresh Capable : No
Border              : No

Neighbor IP address  Generation      State      DR
                    Id               Expires    Refresh    Priority
51.8.8.3             0x53332567      97        No         1
51.8.8.6             0x5332e6d8      94        No         1

Packet Statistics (In/Out)
Hellos               30              15      Bootstraps      10          0
Join/Prunes          12              0      Asserts         0           0
Grafts               0              0      GraftAcks       0           0
State Refresh        0              0

```

The following command displays the elected, active RP for the group 239.255.255.1:

```

show pim rp-set 239.255.255.1
Group      Mask      C-RP      Origin      Priority
224.0.0.0  240.0.0.0  10.10.10.2  Bootstrap  0
224.0.0.0  240.0.0.0  124.124.124.124 Bootstrap  0
224.0.0.0  240.0.0.0  124.124.124.124 static    0
239.255.255.0 255.255.255.0 124.124.124.124 Bootstrap  0
Elected RP is 124.124.124.124

```

The following command displays the PIM configuration for VLAN v3:

```

# show pim v3
PIM SPARSE Interface[2] on VLAN v3 is enabled and up
IP adr: 30.30.30.1 mask: 255.255.255.0 DR of the net: 30.30.30.2
DR Priority      : 1
Passive         : No
Hello Interval   : 30 sec
Neighbor Time out : 105 sec
Join/Prune Interval : 60 sec
Join/Prune holdtime : 210 sec
Trusted Gateway  : none
CRP group List   : pimPolicy with priority 0
Shutdown priority : 1024
Source Specific Multicast : Disabled
State Refresh    : Off
State Refresh Capable : No
Border           : No

Neighbor IP address  Generation      State      DR
                    Id               Expires    Refresh    Priority
30.30.30.2           0x5199b2db      105       No         1

Packet Statistics (In/Out)
Hellos               41              40      Bootstraps      0          20
Join/Prunes          0              0      Asserts         0           0
Grafts               0              0      GraftAcks       0           0
State Refresh        0              0

```

The following is PIM IPv4 show output for the show register policy configuration, including drop counters:

```

sw2.6 # show pim
PIM Enabled, Version 2
PIM CRP Enabled on 1 interfaces
BSR state      : ELECTED ; BSR Hash Mask : 255.255.255.252
Current BSR Info : 61.2.2.2 (Priority 20) expires after 36 sec
Configured BSR Info : 61.2.2.2 (Priority 20) in vlan 11

```

```

CRP Adv Interval      : 60 sec ; CRP Holdtime: 150
BSR Interval         : 60 sec ; BSR Timeout : 130
Cache Timer          : 210 sec ; Prune Timer : 210
Assert Timeout       : 210 sec ; Register Suppression Timeout,Probe: 60, 5
Generation Id        : 0x5331f58c      Dense Neighbor Check : On
PIM-DM State Refresh TTL          : 16
PIM-DM State Refresh Source Active Timer : 210
PIM-DM State Refresh Origination Interval : 60
Threshold for Last Hop Routers: 0 kbps
Threshold for RP                : 0 kbps
Register-Rate-Limit-Interval    : Always active
PIM SSM address range           : None
PIM Register Policy              : swl_rp_filter
PIM Register Policy RP           : None
PIM IP Route Sharing             : Disabled
PIM IP Route Sharing Hash        : Source-Group-Next Hop
Register Checksum to exclude data
Active Sparse Ckts 3 Dense Ckts 0 State Refresh Ckts 0

Global Packet Statistics ( In          Out          Drop )
C-RP-Advs                1153          1155          0
Registers                 3251           0           10
RegisterStops              0           3239          0

```

The following is PIM show output with IP Route Sharing information:

```

sw6.19 # show pim
PIM Enabled, Version 2
PIM CRP Disabled
BSR state              : ACCEPT_PREFERRED ; BSR Hash Mask : 255.255.255.252
Current BSR Info       : 61.2.2.2 (Priority 20) expires after 78 sec
Configured BSR Info    : 0.0.0.0 (Priority 0)
CRP Adv Interval       : 60 sec ; CRP Holdtime: 150
BSR Interval           : 60 sec ; BSR Timeout : 130
Cache Timer            : 210 sec ; Prune Timer : 210
Assert Timeout         : 210 sec ; Register Suppression Timeout,Probe: 60, 5
Generation Id          : 0x5332e6d1      Dense Neighbor Check : On
PIM-DM State Refresh TTL          : 16
PIM-DM State Refresh Source Active Timer : 210
PIM-DM State Refresh Origination Interval : 60
Threshold for Last Hop Routers: 0 kbps
Threshold for RP                : 0 kbps
Register-Rate-Limit-Interval    : Always active
PIM SSM address range           : None
PIM Register Policy              : None
PIM Register Policy RP           : None
PIM IP Route Sharing             : Disabled
PIM IP Route Sharing Hash        : Source-Group-Next Hop
Register Checksum to exclude data
Active Sparse Ckts 11 Dense Ckts 0 State Refresh Ckts 0

Global Packet Statistics ( In          Out          Drop )
C-RP-Advs                0             0             0
Registers                 0           3315            0
RegisterStops             5376          0             0

VLAN      Cid  IP Address      Designated      Flags          Hello J/P  Nbrs
          Cid  IP Address      Router
eight      1  51.8.8.6        / 24 51.8.8.6   rifms-----  30   60    2
eleven     2  51.11.11.6      / 24 51.11.11.6 rifms-----  30   60    0

```

The following shows the output for the `show pim ipv6 v3` command:

```
# show pim ipv6 v3
PIM SPARSE Interface[1] on VLAN v3 is enabled and up
Global IP adr      : 2010::2/64
Local IP adr       : fe80::204:96ff:fe27:f2c6/64
DR of the net      : fe80::204:96ff:fe27:f2c6
DR Priority         : 1

Passive            : No
Hello Interval     : 30 sec
Neighbor Time out  : 105 sec
Join/Prune Interval : 60 sec
Join/Prune holdtime : 210 sec
Trusted Gateway    : none
CRP group List     : none with priority 0
Shutdown priority  : 1024
Source Specific Multicast : Disabled
State Refresh      : Off
State Refresh Capable : No
Border            : No
Secondary Interfaces: 2003::2/ 64
```

address	Generation	State	DR	Neighbor IP
	Id	Expires	Refresh	Priority
fe80::204:96ff:fe26:6c89	0x5192f6f5	101		
No	1			

```
Packet Statistics (In/Out)
  Hellos      5      6
  Bootstraps  0      0
  Join/Prunes 0      0
  Asserts     0      0
  Grafts      0      0
  GraftAcks   0      0
  State Refresh 0      0
```

History

- This command was first available in ExtremeXOS 10.1.
- The PIM-SSM information was added in ExtremeXOS 11.4.
- Border VLAN information was added in ExtremeXOS 12.0.
- The **ipv6** keyword was added to PIM Register Policy Filter feature in ExtremeXOS 15.3.
- DR Priority output was added in ExtremeXOS 15.3.2.
- IP Route Sharing output was added in ExtremeXOS 15.3.2.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

show pim snooping

```
show pim snooping {vlan} vlan_name
```

Description

Displays the PIM snooping configuration for a [VLAN](#).

Syntax Description

<i>vlan_name</i>	Specifies a VLAN.
------------------	-------------------

Default

Disabled.

Usage Guidelines

None.

Example

The following command displays the PIM snooping configuration for the default VLAN:

```
* sw1.79 # show pim snooping "Default"
PIM Snooping          ENABLED
Vlan Default(1)       Snooping DISABLED
```

The following command displays global PIM Snooping configuration:

```
* (pacman debug) sw1.81 # show pim snooping
PIM Snooping          Enabled
<S,G,RPT> Prune       Accept
Vlan                  PIM Snooping    DR          #Nbr
118                   Enabled         NO DR       0
```

History

This command was first available in ExtremeXOS 12.1.

<S,G,RPT> Prune option was added in ExtremeXOS 15.7.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

NEW! show pim tunnel

```
show pim tunnel
```

Description

Displays the PIM configuration on a tunnel.

Syntax Description

tunnel	Specifies PIM tunnel.
---------------	-----------------------

Default

N/A.

Usage Guidelines

Example

The following command displays the global PIM configuration and statistics with tunnel configured:

```
# show pim tunnel mytunnel
PIM SPARSE Interface[2] on VLAN mytunnel is enabled and up
IP adr: 10.0.0.4    mask: 255.255.255.0    DR of the net: 10.0.0.4
DR Priority          : 5
Passive             : No
Hello Interval      : 30 sec
Neighbor Time out   : 105 sec
Join/Prune Interval : 60 sec
Join/Prune holdtime : 210 sec
Trusted Gateway     : none
CRP group List      : none with priority 192
Shutdown priority   : 1024
Source Specific Multicast : Disabled
State Refresh       : Off
State Refresh Capable : No
Border              : No

Neighbor IP address  Generation    State    DR
                    Id              Expires   Refresh  Priority
10.0.0.3             0x62b56975    104      No        1

Packet Statistics (In/Out)
Hellos               122          123    Bootstraps        62          60
Join/Prunes          0           60    Asserts           0           0
Grafts               0           0     GraftAcks         0           0
State Refresh         0           0
```

The following command displays the IPv4 multicast forwarding status on a tunnel (line 8):

```
# show tunnel mytunnel detail
Tunnel Name          : mytunnel
Tunnel Type           : GRE
Flags                 : Up
```

```
Source Address      : 50.1.1.4
Destination Address : 50.1.1.3
IPv4 Forwarding     : Enabled
IPv4 MC Forwarding  : Enabled
IPv6 Forwarding     : Disabled
Virtual Router (VR) : VR-Default
Payload VR          : VR-Default
Interface Address   : 10.0.0.4/24
IP MTU              : 9170
TCP Adjust MSS      : IPv4: Off      IPv6: Off
TCP MSS Value       : IPv4: 0        IPv6: 0
```

History

This command was first available in ExtremeXOS 32.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

show policy

```
show policy {policy-name | detail}
```

Description

Displays the specified policy.

Syntax Description

<i>policy-name</i>	Specifies the policy to display.
detail	Show the policy in detail.

Default

If no policy name is specified, all policies are shown

Usage Guidelines

Use this command to display which clients are using the specified policy. The detail option displays the rules that make up the policy.

Example

The following example displays all policies on the switch:

```
show policy
```

The following is sample output for the command:

```
Switch # sh policy
Policies at Policy Server:
PolicyName                ClientUsage    Client        BindCount
-----
p1                        1             acl           1
p2                        1             acl           1
vlanV1                    1             acl           1
-----
Total Policies : 3
```

History

This command was available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show policy access-list

```
show policy access-list { [list_dot_ruleprofile-index profile_index ]
| [ {matches [app-signature | ether | icmp6type | icmptype |
ipdestsocket | ipfrag | ipproto | ipsourcesocket | iptos | ipttl |
tcpdestportIP | tcpsourceportIP | udpdestportIP | udpsourceportIP ]
{mask mask} {data data} } {actions [ {drop | forward } {cos cos} {-1}
{mirror-destination control_index} {syslog ] } ] } {detail}
```

Description

Displays access list information.

Syntax Description

access-list	Specifies configuring access-list features.
<i>list_dot_rule</i>	Access-list name with optional rule name in format <i>list_name</i> { <i>rule_name</i> }.
profile-index	Specifies the profile index.
<i>profile_index</i>	Defines the profile index (range 1–63). This options shows all access list information associated with the specified profile.
matches	Shows rules with a specific match type, such as match types such as app-signature, ether, etc.
app-signature	Shows application signature specific settings.
ether	Shows type field in Ethernet II packet.
icmp6type	Shows ICMPv6 type.code.
icmptype	Specifies the ICMPv6 type.code.

ipdestsocket	Specifies the destination IP address with optional post-fixed port.
ipfrag	Specifies IP fragmentation flag.
ipproto	Specifies protocol field in IP packet.
ipsourcesocket	Specifies source IP address with optional post-fixed port.
iptos	Specifies IPv4 type of service/IPv6 traffic class field.
ipttl	Specifies IP time to live.
tcpdestportIP	Specifies TCP port destination with optional post-fix IPv4 address.
tcpsourceportIP	Specifies TCP port source with optional post-fix IPv4 address.
udpdestportIP	Specifies UDP port destination with optional post-fix IPv4 address.
udpsourceportIP	Specifies UDP port source with optional post-fix IPv4 address.
mask	Shows rules based on the number of most significant bits to match data value.
<i>mask</i>	Specifies mask value (1-144). Note: You cannot specify "0" because that indicates no mask.
data	Specifies showing rules based on the data (corresponds to type option).
<i>data</i>	Specifies the data value to show (corresponds to type option). You can query for any 'Match data' field of the rule types. The data can be full or partial string or a hexadecimal input that starts with "0x" or "0X" or integer data values (for example: IPTTL, IPTOS, IPProto) Note: Partial matches cannot be found for rule types that have integer values (IPTTL, IPTOS, IPProto, Ether). Since the data field for these rule types only accepts integers (or hex), and are not mixed with IP addresses or ports, it made no sense to do partial matches for these rule types.
actions	Shows rules with a specific action, such as CoS, drop, forward, mirror destination, and Syslog.
drop	Shows rules that are set to drop any packets that match this rule.
forward	Shows rules that are set to forward any packets that match this rule.
-1	Shows rules not assigned a drop or forward action.
cos	Shows rules with the specified Class of Service (CoS).
<i>cos</i>	Specifies the CoS (0-255 or -1).
mirror-destination	Shows rules with the specified mirror destination.
<i>control_index</i>	Specifies the mirror destination control index (1-4).
syslog	Shows rules with Syslog enabled.
detail	Specifies displaying all rule information in detail.

Default

N/A.

Usage Guidelines

This command provides information about all the rules in an access list and the policy profile index that the access list is associated with.

The **detail** option provides detailed information about each rule.



Note

"Rule Hit Count" is cleared whenever the access list is unassigned from a profile, or the profile's assigned access list changes.

Example

The following example shows information for the access-list "ACL1":

```
# show policy access-list list-name ACL1
PID |ACL/Rule/Match |Match Data |Msk|PortStr |ST|TS|VLAN|CoS |Mir|
1 |ACL1
  ace4
    UDPSrcPort |135:192.168.0.1 | 22|
    TCPSrcPort |111:123.190.0.1 | 24|All |NV| |drop| | |
  ace3
    TTL |22 (0x16) | 8|All |NV| | | 3| |
  ace2
    IPTOS |2 (0x2) | 8|All |NV| | | 2| |
  ace1
    Ether |23 (0x17) | 16|All |NV|T |drop| | |
```

Rule Type - Rule Description: Port, MAC Address, IP address etc.

Rule Data - Varies depending on Rule Type

Mask - Mask size for rule data where applicable

ST - V-Volatile NV-NonVolatile

TS - Flags:

T-Traps S-Syslog

For Profile Identifier (PID) Rules:

VLAN - VLAN ID, drop or forward (fwr)

CoS - Class Of Service

Mir - Mirror index if assigned or prohibited (pro)

The following example shows detailed information about rules that are configured to drop packets:

```
# show policy access-list action drop detail
=====
Access-list: :ACL1
Profile Index :1
Rule Name :ace4
Match Type 1 :UDP Source Port
Match Data 1 :135:192.168.0.1
Match Mask 1 :22
Actions
VLAN :0 (Drop)
COS :-1 (Unconfigured)
Mirror :-1 (Unconfigured)
Rule Hit Count : 0
Syslog Status : Disabled
Trap Status : Disabled
Rule Name :ace1
Match Type 1 :Ether Type
Match Data 1 :23
Match Mask 1 :16
Actions
```

```

VLAN          :0      (Drop)
COS           :-1     (Unconfigured)
Mirror        :-1     (Unconfigured)
Rule Hit Count : 222
Syslog Status  : Disabled
Trap Status    : Enabled
=====

```

The following example shows explicit and implicit forward rules:

```

# show policy access-list actions forward
PID |ACL/Rule/Match |Match Data |Msk|PortStr |ST|S|VLAN|CoS |Mir|
31  |ACE
      rule3
        IPDest      |10.4.5.6:22      | 48|
        TCPSrcPort  |62:10.7.8.9      | 48|All |NV|S|fwr| 1| 4|
31  |ACE
      rule4
        TCPDestPort |22               | 16|
        IPProto     |6 (0x6)          | 8|
        Ether       |2048 (0x800)     | 16|All |NV|S|fwr| 7| 2|
31  |ACE
      rule5
        UDPSrcPort  |162:192.1.2.3    | 48|
        UDPDestPort |163:192.3.2.1    | 48|
        TTL         |5 (0x5)          | 8|
        IPTOS       |5 (0x5)          | 8|All |NV|S|fwr| 4| 2|
31  |ACE
      rule7
        IPSource     |10.124.8.9       | 32|
        IPProto      |6 (0x6)          | 8|
        Application  |Health Car ICICIPrude| 72|All |NV|S|fwr| 3| 1|

```

The following example displays implicit CoS rule information:

```

# show policy access-list actions cos -1
PID |ACL/Rule/Match |Match Data |Msk|PortStr |ST|S|VLAN|CoS |Mir|
31  |ACE
      rule1
        IPSource     |10.1.2.3         | 32|
        ICMPType     |8.0              | 16|
        Ether        |2048 (0x800)     | 16|All |NV|S|drop| | |

```

ACL/Rule/Match:

The following example shows partial matches for rules with data "IC":

```

# show policy access-list data IC
PID |ACL/Rule/Match |Match Data |Msk|PortStr |ST|S|VLAN|CoS |Mir|
31  |ACE
      rule7
        IPSource     |10.124.8.9       | 32|
        IPProto      |6 (0x6)          | 8|
        Application  |Health Car ICICIPrude| 72|All |NV|S|fwr| 3| 1|

```

History

This command was first available in ExtremeXOS 30.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show policy access-list action-set

```
show policy access-list action-set {set_id}
```

Description

Shows the pre-defined action set for use in RADIUS Change of Authentication (CoA).

Syntax Description

access-list	Specifies access-list features.
action-set	Defines viewing the set of actions applied to multiple sets of match conditions.
<i>set_id</i>	Specifies viewing a particular action set identified with the global action-set ID (range 1-63).

Default

N/A.

Usage Guidelines

If you do not specify a *set_id*, all actions sets appear.

You can made configuration changes using the `create policy access-list action-set set-id [{drop | forward} {cos cos} {mirror-destination control_index} {syslog}]` command.

Example

The following example shows information for action set "1":

```
# show policy access-list action-set 1
PID |ACL/Rule/Match |Match Data |Msk|PortStr |ST|S|VLAN|CoS |Mir|
| action set 1 | | | | |NV|S| | 3| |
```

History

This command was first available in ExtremeXOS 30.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show policy allowed-type

```
show policy allowed-type ports {detail}
```

Description

Use this command to display a list of currently supported traffic rules applied to the administrative profile for one or more ports.

Syntax Description

allowed-type	Show the list of currently supported and allowed traffic rules applied to the admin rules for each dot1D bridge ports.
<i>ports</i>	Port list.
detail	Displays all information in detail.

Default

If **detail** is not specified, summary information is displayed.

Usage Guidelines

The show policy allowed-type command output displays traffic rule types in attribute ID order (1–31) from left to right. Traffic rule type precedence defaults to the attribute ID order.

The show policy allowed-type command specifies two categories of traffic rule type: supported and allowed. Supported indicates whether the specified port supports the traffic rule type. Allowed is an administrative function. By default, all supported traffic rule types are allowed on the port.

History

This command was first available in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show policy app-signature

```
show policy app-signature
```

Description

Shows minimum time-to-live (TTL) value for Layer 7 policy/application signature.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

The TTL minimum value set by the following command:

```
configure policy app-signature minimum-ttl [none | 1 | 5 | 10]
```

Example

The following example shows the minimum TTL value:

```
# show policy app-signature
Policy Application Signature minimum DNS-reply TTL value is 5 minutes
```

History

This command was first available in ExtremeXOS 30.5.

Limitations

The ExtremeSwitching X435 series switch does not support Layer 7 policy (DNS).

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show policy app-signature group

```
show policy app-signature group {group {name name}} {built-in | custom
    {detail} | detail}
```

Description

Shows policy application signature information.

Syntax Description

app-signature	Shows application signature specific information.
group	Shows application signature group-specific settings.
<i>group</i>	Shows application signature information for the specified group name.
name	Shows application signature display name-specific information.
<i>name</i>	Shows application signature information for the specified display name.
built-in	Shows only built-in signature names (default is both built-in and custom).

custom	Shows only user custom signature names (default is both built-in and custom).
detail	Specifies showing additional signature names and also signature patterns if the name is custom.

Default

By default, both built-in and custom names appear.

Usage Guidelines

Pre-defined/built-in patterns for signature names do not appear.

Example

The following example shows application signature information for the group "E-commerce" and name "Warehouse":

```
# show policy app-signature group "E-commerce" name Warehouse

Application Signature
Index Group Name
-----
5      E-commerce
      Index Type      Signature Name
      -----
      5000 Custom Warehouse
```

The following example shows pre-defined group names:

```
# show policy application group

Application Signature
Index Group Name
-----
1      "Advertising"
2      "Business Applications"
3      "Cloud Computing"
4      "Databases"
5      "E-commerce"
6      "Games"
7      "Peer to Peer"
8      "Protocols"
9      "Search Engines"
10     "Social Networking"
11     "Software Updates"
12     "VPN and Security"
13     "Web Applications"
14     "Web File Sharing"
16     "News and Information"
17     "Storage"
18     "Certificate Validation"
19     "Mail"
20     "Streaming"
21     "Cloud Storage"
22     "Restricted Content"
23     "Corporate Website"
24     "Web Collaboration"
```

```

25  "Real Time and Cloud Communications"
26  "Education"
27  "Health Care"
28  "Travel"
29  "Finance"
30  "Web Content Services"
31  "Sports"
32  "Location Services"

```



Note

Using the **detail** option in the preceding command shows the signature names for each group, as well as the signature patterns for the custom signatures.

The following example shows detailed information for the group "E-Commerce":

```

# show policy application group "E-Commerce" detail
Application Signature
Index Group Name
-----
5      "E-Commerce"
      Index Type      Signature Name
      -----
      1      Built-In 360buy
      2      Built-In Alibaba
      ...
      54     Custom   Warehouse
      Index   Signature Pattern
      -----
      1      bjs.com
      2      costco.com
      3      samsclub.com

```

History

This command was first available in ExtremeXOS 30.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show policy autoclear

```
show policy autoclear interval
```

Description

Shows the interval at which the switch automatically clears rule usage statistics.

Syntax Description

autoclear	Designates viewing information about auto-clearing policy rule usage statistics.
interval	Shows the interval in minutes when the switch automatically clears rule usage. Default is 0 (statistics are not automatically cleared).

Default

By default, the **autoclear** interval is 0, which means that statistics are not automatically cleared.

Usage Guidelines

If you have configured Syslog and/or trap actions to notify you when a policy rule is used by using the following command: `configure policy rule profile_index [{app-signature group group name name} | ether ether | icmp6type icmp6type | icmptype icmptype | ip6dest ip6dest | ipdestsocket ipdestsocket | ipfrag | ipproto ipproto | ipsourcesocket ipsourcesocket | iptos iptos | ipttl ipttl | macdest macdest | macsource macsource | port port | tcpdestportIP tcpdestportIP | tcpsourceportIP tcpsourceportIP | udpdestportIP udpdestportIP | udpsourceportIP udpsourceportIP] {mask mask } {port-string [port_string | all]} {storage-type [non-volatile | volatile]} {drop | forward} {syslog syslog} {trap trap} {cos cos } {mirror-destination control_index} {clear-mirror } , this command shows you the interval when these statistics will be cleared.`

To set this auto-clear interval, use the following command:

```
configure policy autoclear {interval interval}
```

Example

The following example shows the interval for automatically clearing rule usage statistics:

```
# show policy autoclear interval
AutoClear interval is 0 minute(s)
```

History

This command was first available in ExtremeXOS 30.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show policy capability

```
show policy capability
```

Description

This command to display all policy classification capabilities supported by your device.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

The output of this command shows a table listing classifiable traffic attributes and the type of actions, by rule type, that can be executed relative to each attribute. Above the table is a list of all the actions possible on this device.

Example

The following example displays all policy classification capabilities supported by the device:

```
# show policy capability
Priority          Permit          Deny
Precedence Reordering  TCI Overwrite    Rules Table
Rule-Use Accounting  Rule-Use Notification  CoS Table
Longest Prefix Rules  Port Disable Action   Auto Clear Interval
RFC 3580 Map         Policy Enable        Mirror Action
Web Redirect         Authentication Override  NSI
Access List

=====
|          | D |   |   |   |   | F |   |   | D | T |   | Q | |
|          | Y |   |   |   |   | O | S |   | I | C | M | U |
|          | N | A |   |   |   | R | Y |   | S | I | I | A |
|          | A | D | V |   |   | D | W | S | T | A | O | R | R |
|          | M | M | L | C | R | A | L | R | B | V | R | A |
|          | I | I | A | O | O | R | O | A | L | E | O | N |
| SUPPORTED RULE TYPES | C | N | N | S | P | D | G | P | E | R | R | T |
=====
|MAC source address   | X | X |   | X | X | X | X | X | X |   | X |   |
|MAC destination address |   |   |   | X | X | X | X | X | X |   | X |   |
|IPX source address   |   |   |   |   |   |   |   |   |   |   |   |   |
|IPX destination address |   |   |   |   |   |   |   |   |   |   |   |   |
|IPX source socket    |   |   |   |   |   |   |   |   |   |   |   |   |
|IPX destination socket |   |   |   |   |   |   |   |   |   |   |   |   |
|IPX transmission control |   |   |   |   |   |   |   |   |   |   |   |   |
|IPX type field       |   |   |   |   |   |   |   |   |   |   |   |   |
|IPv6 source address  |   |   |   |   |   |   |   |   |   |   |   |   |
|IPv6 destination address |   |   |   | X | X | X | X | X | X |   | X |   |
|IPv6 flow label      |   |   |   |   |   |   |   |   |   |   |   |   |
|IP source address     |   |   |   | X | X | X | X | X | X |   | X |   |
|IP destination address |   |   |   | X | X | X | X | X | X |   | X |   |
|IP fragmentation     |   |   |   | X | X | X | X | X | X |   | X |   |
|UDP port source       |   |   |   | X | X | X | X | X | X |   | X |   |
|UDP port destination  |   |   |   | X | X | X | X | X | X |   | X |   |
|TCP port source       |   |   |   | X | X | X | X | X | X |   | X |   |
|TCP port destination  |   |   |   | X | X | X | X | X | X |   | X |   |
|ICMP packet type      |   |   |   | X | X | X | X | X | X |   | X |   |
|TTL                   |   |   |   | X | X | X | X | X | X |   | X |   |
|IP type of service    |   |   |   | X | X | X | X | X | X |   | X |   |
|IP proto              |   |   |   | X | X | X | X | X | X |   | X |   |
|ICMPv6 packet type    |   |   |   | X | X | X | X | X | X |   | X |   |
|Ether II packet type  |   |   |   | X | X | X | X | X | X |   | X |   |
|LLC DSAP/SSAP/CTRL   |   |   |   |   |   |   |   |   |   |   |   |   |
|VLAN tag              |   |   |   |   |   |   |   |   |   |   |   |   |
|Replace TCI           |   |   |   |   |   |   |   |   |   |   |   |   |
|Application Layer     |   |   |   |   |   |   | X | X | X |   | X |   |
|Access Control List   |   |   |   |   |   |   | X | X | X |   | X |   |
```

Port string	X X	X X X X X X	X	
=====				

History

- This command was first available in ExtremeXOS 16.1.
- The authentication override status was added in ExtremeXOS 22.2
- Mirror action was added in ExtremeXOS 30.2.
- ACL Style Policy support was added in ExtremeXOS 30.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show policy captive-portal

```
show policy captive-portal {web-redirect {redirect_index | all} |  
    listening | rule-use}
```

Description

This command shows the current captive portal configuration for web-redirects, L4 listening ports, and rule use space programming.

Syntax Description

<i>redirect_index</i>	Shows the particular web redirect index indicated (1-10).
all	Shows all web redirect indexes.
listening	Shows captive portal HTTP listening L4 ports.
web-redirect	Shows web redirect information.
rule-use	Shows captive portal rule use.

Default

N/A

Example

The following example shows the currently configured captive portal L4 listening ports:

```
# show policy captive-portal listening  
Captive Portal Listening Ports: 55    33    22
```

The following example shows the currently configured captive portal servers from all web-redirect indexes:

```
# show policy captive-portal web-redirect
Web-redirect Index: 1
  Server Index: 2
    Server Status: enabled
    Server URL:    http://192.168.1.3:5552/path/to/use

Web-redirect Index: 10
  Server Index: 2
    Server Status: disabled
    Server URL:    http://111.222.22.33:802/path/to/use
```

The following example shows all of the captive portal servers at web-redirect index 1 (regardless of configuration):

```
# show policy captive-portal web-redirect 1

Web-redirect Index: 1
  Server Index: 1
    Server Status: disabled
    Server URL:

  Server Index: 2
    Server Status: enabled
    Server URL:    http://192.168.1.3:5552/path/to/use
```

The following example shows the entire captive portal configuration:

```
# show policy captive-portal

Captive Portal Listening Ports: 55    33    22

Web-redirect Index: 1
  Server Index: 1
    Server Status: disabled
    Server URL:

  Server Index: 2
    Server Status: enabled
    Server URL:    http://192.168.1.3:5552/path/to/use

Web-redirect Index: 2
  Server Index: 1
    Server Status: disabled
    Server URL:

  Server Index: 2
    Server Status: disabled
    Server URL:

Web-redirect Index: 3
  Server Index: 1
    Server Status: disabled
    Server URL:

  Server Index: 2
    Server Status: disabled
    Server URL:

Web-redirect Index: 4
  Server Index: 1
    Server Status: disabled
```

```
Server URL:

Server Index: 2
Server Status: disabled
Server URL:

Web-redirect Index: 5
Server Index: 1
Server Status: disabled
Server URL:

Server Index: 2
Server Status: disabled
Server URL:

Web-redirect Index: 6
Server Index: 1
Server Status: disabled
Server URL:

Server Index: 2
Server Status: disabled
Server URL:

Web-redirect Index: 7
Server Index: 1
Server Status: disabled
Server URL:

Server Index: 2
Server Status: disabled
Server URL:

Web-redirect Index: 8
Server Index: 1
Server Status: disabled
Server URL:

Server Index: 2
Server Status: disabled
Server URL:

Web-redirect Index: 9
Server Index: 1
Server Status: disabled
Server URL:

Server Index: 2
Server Status: disabled
Server URL:

Web-redirect Index: 10
Server Index: 1
Server Status: disabled
Server URL:

Server Index: 2
Server Status: disabled
Server URL: http://111.222.22.33:802/path/to/use
```

The following example shows the captive portal rule use status:

```
# show policy captive-portal rule-use
Captive Portal Rule Use: Reserved
```

History

This command was first available in ExtremeXOS 22.3.

The **rule-use** option was added in ExtremeXOS 30.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show policy convergence-endpoint

```
show policy convergence-endpoint
```

Description

This command displays both the Convergence End Points (CEPs) detection state and global policies of all supported CEP types.

Syntax Description

This command has no arguments or variables.

Default

N/A

Example

The following example shows the CEP detection state and global policies of all supported CEP types:

```
# show policy convergence-endpoint

Global Convergence End Point state enabled

Convergence End Point default policies
Type      Policy Index  Policy Name
-----
cisco     12             Cisco IP Phone
lldp-med  0
```

History

This command was first available in ExtremeXOS 22.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show policy convergence-endpoint connections

show policy convergence-endpoint connections **ports** [*port_list* | **all**]

Description

This command displays all learned Convergence End Points (CEPs) information.

Syntax Description

<i>port_list</i>	List of ports to show CEPs on.
all	Specifies that information for CEPs on all ports should appear.

Default

N/A

Example

The following example shows CEPs for port 12:

```
# show policy convergence-endpoint connections ports 12
```

```
Convergence End Point Connection Info for port 12
Endpoint Type      cisco
Policy Index       3
Discovery Time     Mon FEB 06 02:31:42 2008
Firmware Version
Address Type       unknown
Endpoint IP        unavailable
Endpoint MAC       00:04:0d:01:f8:35
```

History

This command was first available in ExtremeXOS 22.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show policy convergence-endpoint ports

show policy convergence-endpoint ports [*port_list* | **all**]

Description

This command displays the enable status of all supported Convergence End Point (CEP) types by port.

Syntax Description

<i>port_list</i>	Shows enable status for the listed ports.
all	Shows enable status for all ports.

Default

N/A

Example

The following example shows the CEP enable status for port 3:

```
# show policy convergence-endpoint ports 3
```

Port	Cisco	LLDP-MED
3	enabled	disabled

History

This command was first available in ExtremeXOS 22.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show policy dynamic

```
show policy dynamic [override | syslog-default | trap-default ]
```

Description

This command displays the status of dynamically assigned policy profile options.

Syntax Description

dynamic	Show status of dynamically assigned profiles.
override	Shows current status of which current profile assignment rules (administrative/dynamic) are overriding the other.
syslog-default	Shows current status of the default that dynamically created rules will have in sending of Syslog messages on rule applied.
trap-default	Shows current status of the default that dynamically created rules will have in sending of SNMP notification on rule applied.

Default

N/A.

Example

The following example shows which current profile admin or dynamic assignment rules are overriding the other:

```
# show policy dynamic override
Administratively assigned rules CURRENTLY OVERRIDE dynamically assigned rules for a given
rule type.
Image    : ExtremeXOS version 30.2.0.31 by release-manager
```

History

This command was first available in ExtremeXOS 16.1.

The Syslog and trap default options were added in ExtremeXOS 30.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show policy invalid

```
show policy invalid [all | {action} {count}]
```

Description

This command displays information about the action the device will apply on an invalid or unknown policy.

Syntax Description

invalid	Show the status of the action the device shall apply on an invalid/unknown policy.
all	Show all the invalid policy status.
action	Specifies the action that the device should take if asked to apply an invalid or unknown policy.
count	The number of times the device has detected an invalid/unknown policy.

Default

N/A.

Example

```
x460-G2-SUM46.44 # show policy invalid all
Current action on invalid/unknown profile is: Apply default policy
Number of invalid/unknown profiles detected: 0
```

History

This command was first available in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show policy mactable

show policy mactable { *vlan_list* }

Description

Use this command to display the VLAN ID - Policy Profile mappings table for all or the specified VLANs.

Syntax Description

mactable	Show VLAN ID - Policy Profile mappings table.
<i>vlan_list</i>	(Optional) VLAN ID or range of IDs (1 to 4094).

Default

If a vlan-list is not specified, mactable entries for all VLANs are displayed.

Example

The following example shows the policy profile mappings table for all VLANs:

```
# show policy mactable
Policy map response      : policy
Policy map last change   : 0 days 12:17:15.27
  VLAN ID                Policy Profile
  1-4094                  0
```

History

This command was first available in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show policy profile

show policy profile {**all** | *profile_index*} {**detail**}

Description

Use this command to display policy profile information.

Syntax Description

<code>profile</code>	Show current policy profile.
<code>all</code>	All profile entries.
<code>profile_index</code>	Profile index.
<code>detail</code>	Displays all information in detail.

Default

If optional parameters are not specified, summary information will be displayed for the specified index or all indexes.

Usage Guidelines

If the rule model is set to ACL Style Policy (access-list), then precedence information does not appear. To change the rule model, use the command `configure policy rule-model [access-list | hierarchical]`.

Example

The following example shows how to display policy information for policy profile 11:

```
show policy profile 11
Profile Index          :11
Profile Name           :
Row Status             :active
Port VID Status        :disabled
Port VID Override      :1
CoS Status             :enabled
CoS                    :7
Syslog on use          :disabled
Trap on use            :disabled
Disable ingress port   :disabled
Replace TCI Status     :enabled
Auth Override Status   :disabled
NSI                    :12345678
Tagged Egress          :
Untagged Egress        :
Forbidden Egress       :
Rule Precedence        :1-2,10,12-19,23,20-22,25,31
                       :MACSource (1), MACDest (2), IPv6Dest (10),
                       :IPSource (12), IPDest (13), IPFrag (14),
                       :UDPSrcPort (15), UDPDestPort (16), TCPSrcPort (17),
                       :TCPDestPort (18), ICMPType (19), ICMP6Type (23),
                       :TTL (20), IPTOS (21), IPProto (22), Ether (25),
                       :Port (31)
Admin Profile Usage     :6
Oper Profile Usage      :6
Dynamic Profile Usage   :none
```

The following example shows all policy profile information:

```
x460-G2-SUM46.52 # show policy profile
|PID|Name|RS|PVID|NSI|CoS|MIR|STDOA|T U F|prec|aSum|dSum|
|1|gear|A|555|none||||||
|2|two|A||none||||||
|5|netadmin|A|10|none|5|||Y||
|10|Employee|A|1024|12345678||||||
```

History

This command was first available in ExtremeXOS 16.1.

The authentication override status was added in ExtremeXOS 22.2

Network Service Identifier (NSI) and ICMPType/ICMP6Type rule precedence information was added in ExtremeXOS 22.5.

Rule precedence order change supported in ExtremeXOS 32.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show policy resource-profile

```
show policy resource-profile {[default | less-acl [more-ipv4 |more-ipv4-
no-ipv6 | more-ipv4-no-l2 |more-ipv4-no-mac-no-ipv6 | more-ipv4-no-
mac-no-ipv6-no-l2] |more-ipv4-no-ipv6 | more-ipv4-no-mac-no-ipv6 |
more-mac-no-ipv6] {profile-modifier [ {no-mac} {no-ipv4} {no-ipv6}
{no-l2} ] ] }
```

Description

Shows policy resource settings and usage for currently set profile, or optionally shows the potential resource setting for a specified, different profile.

Syntax Description

default	Shows potential policy resource settings if configuring default profile.
less-acl	Shows potential policy resource settings if configuring profile that removes some access list resources to be used for rules.
more-ipv4	Shows potential policy resource settings if configuring profile that adds IPv4 rules.
more-ipv4-no-ipv6	Shows potential policy resource settings if configuring profile that adds IPv4 rules and removes IPv6 rules.
more-ipv4-no-l2	Shows potential policy resource settings if configuring profile that adds IPv4 rules and removes L2 rules.
more-mac-no-ipv6	Shows potential policy resource settings if configuring profile that adds MAC rules and removes IPv6 rules.

more-ipv4-no-mac-no-ipv6-no-l2	Shows potential policy resource settings if configuring profile that adds IPv4 rules and removes MAC, IPv6, and L2 rules.
more-ipv4-no-mac-no-ipv6	Shows potential policy resource settings if configuring profile that adds IPv4 rules and removes MAC and IPv6 rules.
profile-modifier	Shows potential policy resource settings if configuring a profile-modifier.
no-mac	Shows potential policy resource settings if configuring a profile-modifier that removes all MAC rules.
no-ipv4	Shows potential policy resource settings if configuring a profile-modifier that removes all IPv4 rules.
no-ipv6	Shows potential policy resource settings if configuring a profile-modifier that removes all IPv6 rules.
no-l2	Shows potential policy resource settings if configuring a profile-modifier that removes all L2 rules.

Default

None.

Usage Guidelines

Pressing **ENTER** rather than selecting any options for the command shows the current resource profile settings.

Example

The following example shows the current resource profile settings:

```
# show policy resource-profile

Current Configured Profile: default
Current Profile Modifier   : no-l2

      MAC  IPv6  IPv4   L2
      Rules Rules Rules Rules
      ----
Max    512   512   440    0
Used    0     0     0     0

L2 ether rules are accounted for in the IPv4 group
```

The following example shows what resources would look like when using a different profile on a stack (with potentially different slot types):

```
show policy resource-profile more-mac-no-ipv6

      MAC  IPv6  IPv4   L2
Slot Rules Rules Rules Rules Type
-----
1      512    0   256   184 X450G2-24t-G4
2      512    0   256   184 X450G2-48t-10G4
3      512    0   256   184 X450G2-48t-G4
4        0    0   256   184 X620-10x
Max      0    0   256   184 (Stack)
```

```
Used      3      1      5      0 (Stack)
```

```
Max : This row shows the maximum resources available for the stack
      if the specified profile is chosen.
```

The following example shows what resources would look like when using a different profile on a single switch:

```
show policy resource-profile more-mac-no-ipv6
```

	MAC	IPv6	IPv4	L2	
Slot	Rules	Rules	Rules	Rules	Type
1	0	0	256	184	X620-10x
Max	0	0	256	184	(Switch)
Used	1	0	3	0	(Switch)

```
Max : This row shows the maximum resources available for the switch
      if the specified profile is chosen.
```

History

This command was first available in ExtremeXOS 22.1.

Profile modification information was added in ExtremeXOS 22.4.

Additional profiles (no-l2) and profile modifier were added in ExtremeXOS 30.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show policy rule

```
show policy rule {all | app-signature | {profile-index profile_index
| admin-profile} ether {ether} | icmp6type {icmp6type} | icmp6type
{icmp6type} | ip6dest {ip6dest} | ipdest {ipdest} | ipfrag |
ipproto {ipproto} | ipsource { ipsource } | iptos { iptos }
| ipttl { ipttl } | macdest { macdest } | macsource
{ macsource } | port { port } | tcpdestportIP { tcpdestportIP } |
tcpsourceportIP { tcpsourceportIP } | udpdestportIP { udpdestportIP }
| udpsourceportIP { udpsourceportIP }} {mask mask } {port-string
[ port_string | all]} {storage-type [non-volatile | volatile]} {drop
| forward} {cos cos | admin-pid admin_pid }} {detail | wide}
```

Description

Use this command to display policy classification and admin rule information.

Syntax Description

rule	Show current Policy Rule.
all	Optional, show all policy rules
app-signature	Specifies application signature specific settings.
profile-index	Optional: Specify the profile index
admin-profile	Optional: Show rule based on Policy ID of 0
mask	Optional: Show rule based on the number of most significant bits to match data value.
<i>mask</i>	Optional: Show rule based on the number of most significant bits to match data value. Range = 1-144.
port-string	Optional: Show rule based on the port number on which this rule is applied; single port in port-string format.
<i>port-string</i>	Optional: Show rule based on the port number on which this rule is applied; single port in port-string format.
storage-type	Optional: Show rule based on its non-volatile storage type (V - volatile; NV - non-volatile).
non-volatile	Show rule with non-volatile storage type.
volatile	Show rule with volatile storage type.
drop	Show rules that are set to 'drop' any packets which match this rule.
forward	Show rules that are set to 'forward' any packets which match this rule.
cos	Optional: Show rules with Class of Service.
<i>cos</i>	Optional: Show rules with Class of Service (0-255) or -1.
admin-pid	Policy ID.
<i>admin-pid</i>	Policy ID. Range = 0 - 102.
wide	Optional: Extend the concise view beyond 80 columns to display complete rule data.
detail	Optional: show all rule information in detail.
port	Port string.
<i>port</i>	Port string - (data: 1; mask: 16).
macdest	MAC destination address.
<i>macdest</i>	MAC destination address - (data: a-b-c-d-e-f; mask: 1-48).
ip6dest	IPv6 address.
<i>ip6dest</i>	IPv6 address (data: aaaa::bbbb; mask 1-128).
ipsource	Source IP address.
<i>ipsource</i>	Source IP address - (data: a.b.c.d; mask: 1-32).
ipdest	Destination IP address.
<i>ipdest</i>	Destination IP address - (data: a.b.c.d.; mask: 1-32).
ipfrag	IP fragmentation flag.
tcpdestportIP	TCP port dst with optional post-fix IPv4 address.

<i>tcpdestportIP</i>	TCP port dst with optional post-fix IPv4 address - (data: ab[:c.d.e.f]); mask: 1-48.
udpdestportIP	UDP port dst with optional post-fix IPv4 address.
<i>udpdestportIP</i>	UDP port dst with optional post-fix IPv4 address - (data: ab[:c.d.e.f]); mask: 1-48.
tcpsourceportIP	TCP port src with optional post-fix IPv4 address.
<i>tcpsourceportIP</i>	TCP port src with optional post-fix IPv4 address - (data: ab[:c.d.e.f]); mask: 1-48.
udpsourceportIP	UDP port src with optional post-fix IPv4 address.
<i>udpsourceportIP</i>	UDP port src with optional post-fix IPv4 address - (data: ab[:c.d.e.f]); mask: 1-48.
ipttl	IP time to live.
<i>ipttl</i>	IP time to live - (data: 0-255).
iptos	IPv4 type of service / IPv6 traffic class field.
<i>iptos</i>	IPv4 type of service / IPv6 traffic class field - (data: 0-255; mask: 1-8).
ipproto	Protocol field in IP packet.
<i>ipproto</i>	Protocol field in IP packet - (data: 0-255 or 0-0xFF; mask: 1-8).
ether	Type field in Ethernet II packet.
<i>ether</i>	Type field in Ethernet II packet - (data: 0-65535 or 0x0-0xFFFF; mask: 1-16).
icmp6type	Specifies type code in ICMPv6 packet.
<i>icmp6type</i>	ICMPv6 type code [(data: 123.456 (dotted-decimal) or AB-CD (dashed-hexadecimal))] mask: 1-16).
icmptype	Specifies type code in ICMP packet.
<i>icmptype</i>	ICMP type code (data: a.b; mask: 1-16).

Default

- If port-string, cos and storage-type are not specified, all rules related to other specifications will be displayed.
- If -verbose is not specified, summary information will be displayed.
- If -wide is not specified, an 80 character display width is used.

Usage Guidelines

Use this command to display policy classification and admin rule information.

Example

The following example shows policy classification and admin rule information:

```
# show policy rule
Admn|Rule Type   |Rule Data                |Msk|PortStr  |RS|ST|dPID|aPID|Mir|
adm|MACSource    |00-77-77-77-00-20       |48|1        |A|V|5|   |   |
```



```

adm|MACSource |00-77-77-77-00-21 | 48|4 | A| V| 5| | |
adm|Port |1 | 16|1 | A|NV| | 22| |
adm|Port |4 | 16|4 | A|NV| | 22| |
PID |Rule Type |Rule Data |Msk|PortStr |RS|ST|VLAN|CoS |Mir|
5 |Ether |2048 (0x800) | 16|All | A|NV|fwr| | 1|
5 |Ether |33079 (0x8137) | 16|All | A|NV|fwr| | 1|

Rule Type - Rule Description: Port, MAC Address, IP address etc.
Rule Data - Varies depending on Rule Type
Mask - Mask size for rule data where applicable
RS - RowStatus:
    A-Active NS-NotInService NR-NotReady CG-CreateAndGo CW-CreateAndWait D-Destroy
ST - V-Volatile NV-NonVolatile
For Admin Profile Rules (Admn):
    dPID - Dynamic Profile Index
    aPID - Admin Profile Index
For Profile Identifier (PID) Rules:
    VLAN - VLAN ID, drop or forward (fwr)
    CoS - Class Of Service
Mir - Mirror index if assigned

```

The following example shows detailed policy classification and admin rule information:

```

# show policy rule detail
=====
Profile Index      :Admin-Profile
Rule Type          :Port string
Rule Data          :26
Mask               :16
Port               :26
-----
Status             :active
Storage Type       :nonVolatile
Operational-PID    :-1
Admin-PID          :1
=====
Profile Index      :1
Rule Type          :MAC source address
Rule Data          :00-00-00-00-00-10
Mask               :48
Port               :All ports
-----
Status             :active
Storage Type       :nonVolatile
VLAN               :-1 (Unconfigured)
COS                :-1 (Unconfigured)
Mirror             :0 (Prohibited)

Rule Hit Count     : 0
Audit Syslog Status : Prohibit
Audit Trap Status  : Prohibit
=====
Profile Index      :1
Rule Type          :Port string
Match Type 1       :MAC source address
Match Data 1       :192.168.123.100
Match Mask 1       :32
Match Type 2       :IP source address
Match Data 2       :00-00-00-00-00-10
Match Mask 2       :48
Port               :All ports
-----
Status             :active

```

```

Storage Type      :nonVolatile
VLAN              :0      (Drop)
COS               :-1     (Unconfigured)
Mirror            :0      (Prohibited)

Rule Hit Count    : 0
Audit Syslog Status : Enabled
Audit Trap Status  : Prohibit
=====

```

History

This command was first available in ExtremeXOS release 16.1.

ICMP and ICMPv6 type information added in ExtremeXOS 22.5.

Mirror information and rule usage counter information were added in ExtremeXOS 30.2.

The **app-signature** option was added in ExtremeXOS 30.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show policy rule port-hit

```
show policy rule port-hit {data} {detail} {wide}
```

Description

Shows a list of used rules when Syslog/trap actions are configured for policy.

Syntax Description

port-hit	Shows ports matching the rules.
data	Shows rule based on the data (corresponds to type option).
detail	Shows all rule information in detail.
wide	Extends the concise view beyond 80 columns to display complete rule data.

Default

N/A.

Usage Guidelines

If you have configured Syslog and/or trap actions to notify you when a policy rule is used by using the following command: `configure policy rule profile_index [{app-signature group group name name} | ether ether | icmp6type icmp6type`

```
| icmp-type icmp-type | ip6dest ip6dest | ipdestsocket ipdestsocket |
ipfrag | ipproto ipproto | ipsourcesocket ipsourcesocket | iptos
iptos | ipttl ipttl | macdest macdest | macsource macsource | port
port | tcpdestportIP tcpdestportIP | tcpsourceportIP tcpsourceportIP |
udpdestportIP udpdestportIP | udpsourceportIP udpsourceportIP ] {mask
mask } {port-string [ port_string | all]} {storage-type [non-volatile
| volatile]} {drop | forward} {syslog syslog} {trap trap} {cos cos }
{mirror-destination control_index} {clear-mirror} , this command shows you
information about the rules that have been used.
```

You can clear this information by using the command `clear counters policy`.

Example

The following example shows rule usage information:

```
# show policy rule port-hit
PID |Rule Type |Rule Data |Msk|PortStr |RS|ST|TS|VLAN|CoS |Mir|
1 |MACSource |00-00-77-77-00-01 | 48|25 | A|NV|TS|fwr| |None|
1 |MACSource |00-00-77-77-00-02 | 48|All | A|NV|T |fwr| |None|
```

The following example shows detailed rule usage information:

```
# show policy rule port-hit detail
=====
Profile Index      :1
Rule Type          :MAC source address
Rule Data          :00-00-77-77-00-01
Mask               :48
Port               :25
- - - - -
Status             :active
Storage Type       :nonVolatile
VLAN               :4095 (Forward)
COS                :-1 (Unconfigured)
Mirror             :None

Rule Hit Count     : 429
Audit Syslog Status : Enabled
Audit Trap Status  : Enabled
=====
Profile Index      :1
Rule Type          :MAC source address
Rule Data          :00-00-77-77-00-02
Mask               :48
Port               :All ports
- - - - -
Status             :active
Storage Type       :nonVolatile
VLAN               :4095 (Forward)
COS                :-1 (Unconfigured)
Mirror             :None

Rule Hit Count     : 410
Audit Syslog Status : Prohibit
Audit Trap Status  : Enabled
=====
```

History

This command was first available in ExtremeXOS 30.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show policy slices

show policy **slices**

Description

Shows the existing usage of policy slices.

Syntax Description

slices	Shows look-up stage TCAM resources.
---------------	-------------------------------------

Default

N/A.

Example

The following example shows the existing usage of policy slices:

```
# show policy slices
Configured # tci-overwrite slices : 2
Maximum # slices                  : 4
Current Configured shared slices           : 2
Current Configured guaranteed DynAcl percentage: 40%
Current Configured guaranteed L7 percentage  : 40%
```

History

This command was first available in ExtremeXOS 30.4.

Information about configured guaranteed Layer 7 policy and dynamic ACL percentages was added in ExtremeXOS 30.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show policy state

show policy state

Description

This command shows the current policy state.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Example

```
# show policy state
Policy is currently: DISABLED
Rule-model          : access-list
```

History

This command was first available in ExtremeXOS 16.1.

ACL Style Policy status was added in ExtremeXOS 30.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show policy syslog

```
show policy syslog [machine-readable] [extended-format] [every-time]
```

Description

Shows the Syslog parameters for policy rules.

Syntax Description

syslog	Designates showing the Syslog parameters for policy rules.
machine-readable	Shows whether Syslog messages are shown in decimal or hexadecimal format.
extended-format	Shows whether extended formatting of rule usage messages is enabled or disabled.
every-time	Shows whether Syslog messages are sent every time a rule is used or just the first time.

Default

N/A.

Usage Guidelines

None.

Example

The following example shows the Syslog parameters for policy rules:

```
# show policy syslog
Syslog messages will be logged in DECIMAL format.
Extended format DISABLED on rule usage messages.
Syslog messages sent on EACH USE of rule.
```

History

This command was first available in ExtremeXOS 30.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show policy vlanauthorization

```
show policy vlanauthorization {port port_list }
```

Description

This command displays VLAN Authorization information for a port or list of ports.

Syntax Description

port	Show VLAN Authorization information for a specific set of ports.
port_list	Specifies the list of ports to show VLAN Authorization information for.

Default

If no parameters are specified, all VLAN Authorization configuration information is displayed.

Usage Guidelines

None.

Example

This example shows how to display VLAN Authorization configuration information for ports 1:1-3:

```
X450G2-48t-10G4.3 # show policy vlanauthorization port 1-3
VLAN Authorization Global Status: ENABLED
      Admin      Oper
Port  Status    Egress  Egress  VLAN ID
=====
1      enabled   untagged untagged none
2      enabled   untagged untagged none
3      enabled   untagged untagged none
```

History

This command was first available in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show ports

```
show ports {port_list | tag tag} {no-refresh | refresh}
```

Description

Displays port summary statistics.

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
tag	Specifies an 802.1Q or 802.1ad tag value.
no-refresh	Specifies a static snapshot of the data.
refresh	Specifies a continuous refresh of output.

Default

N/A.

Usage Guidelines

Use this command to display the port number, display string, and some of the port states in tabular form.

The VLAN name is displayed only if that port contains a single VLAN. If the port contains more than one VLAN, then the number of VLANs is displayed.

The tag value may be associated with either a VMAN or a VLAN.

The default display string for Integrated Application Hosting (IAH) dedicated ports on ExtremeSwitching X465 series switches is “Insight” to indicate that these are to be used for virtualization. On X465-24MU and and X465-24XE switches, the dedicated ports are 33-34; on X465-24MU-24W, and X465i-48W switches, the dedicated ports are 57-58. You can change this display string (configure ports *port_list* **display-string** *string*). If you unconfigure the display string, it is reset to “Insight”.

Example

The following command displays information on ports 2, 3, and 12 on a switch:

```
show ports 2,3,12
Port Summary Monitor
Thu Feb 14 14:19:50 2008
Port  Display          VLAN Name      Port  Link  Speed  Duplex
#      String            (or # VLANs)  State State Actual Actual
=====
2    2nd-Floor-Lab    Lab-Backbone   E     A     1000   FULL
3          Building2      E     A D
12 AllBackboneLANs (34) E     R               FULL
=====
Port State: D-Disabled, E-Enabled, F-Disabled by link-flap detection,
            L-Disabled due to licensing, M-MKA enabled
Link State: A-Active, R-Ready, NP-Port Not Present, L-Loopback,
            D-ELSM enabled but not up,
            d-Ethernet OAM enabled but not up,
            B-MACsec enabled but blocked awaiting authentication
U->page up  D->page down ESC->exit
```

Restricted optics will show similarly to unsupported optics (!) in the show ports command by use of a '\$'. Additionally, when a 3rd-party license has been installed, previously restricted optics are shown using a '%', because those optics are un-restricted by means of the license. Here is an example output:

```
show port conf
Port Configuration Monitor
Wed Sep 5 20:35:54 2012
Port  Virtual  Port  Link  Auto  Speed  Duplex  Flow  Load  Media
      router   State State Neg  Cfg Actual Cfg Actual Cntrl Master Pri Red
=====
1    VR-Default E     R    OFF 40000      FULL      $Q+SR4
2    VR-Default E     NP  OFF 10000      FULL      NONE
3    VR-Default E     NP  OFF 10000      FULL      NONE
4    VR-Default E     NP  OFF 10000      FULL      NONE
5    VR-Default E     R    OFF 40000      FULL      NONE
6    VR-Default E     NP  OFF 10000      FULL      NONE
7    VR-Default E     NP  OFF 10000      FULL      NONE
=====
indicates Port
> indicates Port Display Name truncated past 8 characters
Link State: A-Active, R-Ready, NP-Port Not Present, L-Loopback type
Port State: D-Disabled, E-Enabled, L-License Disabled
Media: !-Unsupported, $-Unlicensed detail" for redundant media type
Media Red: * - use "show port info detail" for redundant media type
Flow Cntrl: Shows link partner's abilities. NONE if Auto Neg is OFF
0->Clear Counters U->page up D->page down ESC->exit
```

History

This command was first available in ExtremeXOS 12.1.

The tag value was added in ExtremeXOS 12.4.4.

Show output for 3rd party restricted optics support was added in ExtremeXOS 15.3.

The **refresh** keyword was added in ExtremeXOS 16.1.

Flag to show ports disabled by excessive link-flapping added in v 22.1.

Flag (L) to show port disabled due to lack of port speed license (applies to ExtremeSwitching X870-96x-8c switches only).

Flag (B) to show port blocked by MAC Security was added in ExtremeXOS 30.1.

"Insight" display string appears for IAH dedicated ports of ExtremeSwitching X465-24MU and X465-24MU-24W switches for ExtremeXOS 30.2.

"Insight" display string appears for IAH dedicated ports of ExtremeSwitching X465-24XE and X465i-48W switches for ExtremeXOS 30.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show ports advertised

```
show ports {port_list} advertised
```

Description

Shows the remote end advertised ability.

Syntax Description

ports	Specifies ports.
<i>port_list</i>	Lists the ports that you want remote end advertised ability information on.
advertised	Selects showing auto-negotiation advertised capabilities.

Default

N/A.

Usage Guidelines

Auto-negotiation is an Ethernet feature that facilitates the selection of port speed, duplex, and flow control between the two members of a link. It first shares these capabilities, and then selects the fastest transmission mode that both ends of the link can support.

This command shows the remote end of the link's advertised ability, up to 1Gbps speed only, for duplex and flow control. This command is not supported on Extended Edge Switching bridge port extenders (BPE) ports.

Example

The following example shows the advertised ability for the remote end for ports 1 and 9:

```
# show port 1 9 advertised
```

			10BaseT		100BaseT		100X		1000BaseT		1000X		
Pause													
Port	Half	Full	Half	Full	Half	Full	Half	Full	Half	Full	TX	RX	Asym
*1	CL-	CL-	CL-	CL-	---	---	---	---	---	---	CLR	CLR	---
!9	CL-	CL-	CL-	CL-	---	---	---	---	---	---	C--	C--	---

History

This command was first available in ExtremeXOS 30.2.

Platform Availability

VLAN

show ports anomaly

```
show ports {port_list | tag tag} anomaly {no-refresh | refresh}
```

Description

Displays statistics of anomaly violation events in real time.

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
<i>tag</i>	Specifies an 802.1Q or 802.1ad tag value.
no-refresh	Specifies a static snapshot of data.
refresh	Continous refresh of output.

Default

N/A.

Usage Guidelines

If you do not specify a port number, range of ports, or tag value, statistics are displayed for all ports. To clear the counters, use the `clear counters ports` command. The default display is a constantly

refreshing real-time display. If you specify the no-refresh parameter, the system displays a snapshot of the data at the time you issue the command.

This command takes effect after enabling anomaly-protection.

The tag value may be associated with either a VMAN or a VLAN.

Example

The following command displays real-time anomaly statistics on a switch:

```
show ports anomaly
```

Following is sample output from this command:

```
Port Statistics Thu Nov  9 22:44:31 2006
Port Link      Rx Pkt ===== Anomaly Violation =====
State Count L3 Count      L4 Count      ICMP Count      Frag Count
=====
1 A 191585 1 2 0 0
2 R 0 0 0 0 0
3 R 0 0 0 0 0
4 R 0 0 0 0 0
5 R 0 0 0 0 0
6 R 0 0 0 0 0
7 R 0 0 0 0 0
8 R 0 0 0 0 0
9 R 0 0 0 0 0
10 R 0 0 0 0 0
11 R 0 0 0 0 0
12 A 178024 0 0 0 0
13 A 196956 0 0 0 0
14 R 0 0 0 0 0
15 R 0 0 0 0 0
16 R 0 0 0 0 0
17 R 0 0 0 0 0
=====
Link State: A-Active, R-Ready, NP-Port Not Present, L-Loopback
0->Clear Counters  U->page up  D->page down  ESC->exit
```

History

This command was first available in ExtremeXOS 12.0.

The tag value was added in ExtremeXOS 12.4.4.

The **refresh** keyword was added in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show ports buffer

show **ports** *port_list* **buffer**

Description

Displays the packet buffer organization for the specified ports.

Syntax Description

<i>port_list</i>	Optionally specifies the list of ports, or slots and ports, for which packet buffer information is displayed. If the <i>port_list</i> is omitted then packet buffer information is displayed for all ports in the system.
------------------	---

Default

N/A.

Usage Guidelines

This command shows the packet buffer organization for the specified ports.

The *port_list* can span multiple ranges. The packet buffer description for each such port range is displayed.

Since ports and packet buffer are grouped by the hardware, the command displays the range of ports that share the same packet buffer.

The Total Packet Buffer Size for the port range is displayed in bytes, along with an indication of whether or not the user has configured over-commitment of the packet buffer (not overcommitted by default).

The amount of Reserved Buffer allocated to each port and QoS Profile is shown for the ports in the user-specified *port_list*. To configure the reserved buffer, use the `configure qosprofile qosprofile maxbuffer percentage ports port_list` command.

The Total Shared Buffer Size displayed is the Total Packet Buffer Size minus the total Reserved Buffer allocated to all ports and QoS profiles in the port range. Note that some packet buffer is also reserved to internal ports.

For each port, the maximum of the Total Shared Buffer Size that the port is allowed to use (Max Shared Buffer Usage) is shown both as an absolute number of bytes and as a percentage of the Total Shared Buffer Size. A port's Max Shared Buffer Usage may be configured using the command `configure ports {port_list} shared-packet-bufferpercentage`

Note the configured percentage may be different than the displayed percentage. This is because more recent hardware can only allocate shared packet buffer in steps, while older hardware can precisely allocate the requested percentage.

The more recent hardware dynamically adjusts each port's shared buffer usage limit based on simultaneous usage by multiple ports and QoS profiles, automatically providing fair usage of the

shared buffer among the ports and QoS profiles that are currently demanding buffer space. This allows larger packet buffer usage bursts on a port when other ports are not using shared buffer. This dynamic adjustment cannot be observed with this command since only the maximum possible limits are displayed.

The VLAN name is displayed only if that port contains a single VLAN. If the port contains more than one VLAN, then the number of VLANs is displayed.

The tag value may be associated with either a VMAN or a VLAN.

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show ports collisions

```
show ports {mgmt | port_list | tag tag} collisions {no-refresh |  
refresh }
```

Description

Displays real-time collision statistics.

Syntax Description

mgmt	Specifies the management port.
<i>port_list</i>	Specifies one or more ports or slots and ports.
<i>tag</i>	Specifies an 802.1Q or 802.1ad tag value.
no-refresh	Specifies a static snapshot of data.
refresh	Continuous refresh of output.

Default

Real-time statistics.

Usage Guidelines

If you do not specify a port number, range of ports, or tag value, statistics are displayed for all ports. To clear the counters, use the `clear counters ports` command. The default display is a constantly refreshing real-time display. If you specify the no-refresh parameter, the system displays a snapshot of the data at the time you issue the command.

The tag value may be associated with either a VMAN or a VLAN.

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 ports 1 and 2 on a switch:

```
show ports 1-2 collisions
```

Following is sample output from this command:

Port Collision Monitor																
Port	Link	Collision Histogram														
State	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
1	A	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
2	R	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Link State: A-Active R-Ready, NP-Port not present, L-Loopback																

The numbers 1 to 16 represent the number of collisions encountered prior to successfully transmitting the packet; this is applicable only for half-duplex links.

History

- This command was first available in ExtremeXOS 10.1.
- The no-refresh variable was added in ExtremeXOS 11.3.
- The tag value was added in ExtremeXOS 12.4.4.
- The **refresh** keyword was added in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show ports configuration

```
show ports {mgmt | port_list | tag tag} configuration {no-refresh | refresh}
```

Description

Displays port configuration statistics, in real time or snapshot.

Syntax Description

mgmt	Specifies the management port.
<i>port_list</i>	Specifies one or more ports or slots and ports.

<i>tag</i>	Specifies an 802.1Q or 802.1ad tag value.
no-refresh	Specifies a static snapshot of data.
refresh	Specifies a continuous refresh of output.

Default

Real-time statistics.

Usage Guidelines

If you do not specify a port number, range of ports, or tag value, configuration statistics are displayed for all ports. If you specify the no-refresh parameter, the system displays a snapshot of the data at the time you issue the command.

The tag value may be associated with either a VMAN or a VLAN.

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

This command displays port configuration, which includes:

- Virtual router.
- Port state.
- Link state.
- Autonegotiation information.
- Link speed.
- Duplex mode.
- Flow control.
- Load sharing information.
- Link media information.



Note

On 10 Gbps ports, the Media Primary column displays NONE when no module is installed, and SR, LR, or ER depending on the module installed when there is one present. Combination ports display Autonegotiation, Link speed, and Duplex mode information only for the current primary medium.

Example

The following command displays the port configuration for all ports on a switch:

```
# show port configuration no-refresh
Port Configuration
Port      Virtual      Port  Link  Auto   Speed      Duplex   Flow   Load   Media
          router      State State Neg   Cfg Actual Cfg Actual Cntrl Master Pri Red
=====
1         VR-Default E      R    OFF 10000      FULL
2         VR-Default E      R    OFF 10000      FULL
3         VR-Default E      R    OFF 1000       FULL
4         VR-Default E      R    OFF 10000      FULL
                                     #SF+_SR
```

```

5      VR-Default E      R  OFF 10000      FULL      NONE
6      VR-Default E      A  ON  1000 1000    FULL FULL  NONE    1000T
7      VR-Default E      A  ON  1000 1000    FULL FULL  NONE    1000T
8      VR-Default E      R  OFF 10000      FULL      NONE
9      VR-Default E      R  OFF 1000      FULL      #NONE
10     VR-Default E      R  OFF 10000      FULL      #NONE
=====

```

```

> indicates Port Display Name truncated past 8 characters
Link State: A-Active, R-Ready, NP-Port Not Present, L-Loopback
Port State: D-Disabled, E-Enabled, L-License Disabled
Media: !-Unsupported, $-Unlicensed, #-LRM/MACsec Adapter
Media Red: * - use "show port info detail" for redundant media type
Flow Cntrl: Shows link partner's abilities. NONE if Auto Neg is OFF

```

Example

The following command displays the indication that a 25G lane based optic has been inserted into port partitioned for 10G lanes:

```

(pacman debug) 5520-24X-EXOS.7 # show port 25 config no
Port Configuration
Port      Virtual      Port Link Auto   Speed      Duplex      Flow Load Media
          router      State Neg  Cfg Actual  Cfg Actual Cntrl Master Pri Red
=====
25      VR-Default E      R  OFF 40000      FULL      ?Q28+SWDM4
=====
> indicates Port Display Name truncated past 8 characters
Link State: A-Active, R-Ready, NP-Port Not Present, L-Loopback
Port State: D-Disabled, E-Enabled, L-License Disabled
Media: !-Unsupported, $-Unlicensed, #-LRM/MACsec Adapter, ?-Mismatches Partition
Media Red: * - use "show port info detail" for redundant media type
Flow Cntrl: Shows link partner's abilities. NONE if Auto Neg is OFF
(pacman debug) 5520-24X-EXOS.8 #

```

Example

The following command displays the indication that a 100G optic has been inserted into an unlicensed port:

```

(pacman debug) 5520-24X-EXOS.7 # show port 1 config no
Port Configuration
Port      Virtual      Port Link Auto   Speed      Duplex      Flow Load Media
          router      State Neg  Cfg Actual  Cfg Actual Cntrl Master Pri Red
=====
25      VR-Default E      R  OFF 10000      FULL      ?$Q28+SWDM4
=====
> indicates Port Display Name truncated past 8 characters
Link State: A-Active, R-Ready, NP-Port Not Present, L-Loopback
Port State: D-Disabled, E-Enabled, L-License Disabled
Media: !-Unsupported, $-Unlicensed, #-LRM/MACsec Adapter, ?-Mismatches Partition
Media Red: * - use "show port info detail" for redundant media type
Flow Cntrl: Shows link partner's abilities. NONE if Auto Neg is OFF

```

History

This command was first available in ExtremeXOS 10.1.

The Port not present and Media variables were added in ExtremeXOS 11.2.

The no-refresh variable was added in ExtremeXOS 11.3.

The tag value was added in ExtremeXOS 12.4.4.

This command was updated to include flags that indicate the summary status of a non-Extreme optical device in ExtremeXOS 15.4.

The **refresh** keyword was added in ExtremeXOS 16.1.

Flag (L) to show port disabled due to lack of port speed license (applies to ExtremeSwitching X870-96x-8c switches only).

Flag (#) to show that LRM/MACsec adapter is connected to a port was added in ExtremeXOS 30.1.

Mismatched port compatibility indication was added in ExtremeXOS 31.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show ports congestion

```
show ports port_list congestion {no-refresh | refresh} {port-number}
```

Description

Displays the port egress congestion statistics (dropped packets) for the specified ports on the front panel.

Syntax Description

<i>port_list</i>	Specifies one or more slots and ports.
no-refresh	Specifies a static snapshot of data instead of the default dynamic display.
refresh	Specifies a continuous refresh of output.
port-number	Displays port number instead of display string.

Default

Displays the port congestion statistics for all ports in real-time.

Usage Guidelines

The bottom line in the real-time display shows keys that you can press to change the display. For example, you can clear the counters or page up or down through the list of ports.



Note

If you are displaying congestion statistics in real time and another CLI session resets the counters for a port you are monitoring, the counters displayed in your session for that port are also reset.

If you specify the no-refresh parameter, the system displays a snapshot of the data at the time you issue the command.



Note

Packets can be dropped at multiple locations along the path through the hardware. The per-port congestion counters count all dropped packets for all ports.

If you do not specify a port number or range in the command, dropped packet counts are displayed for all ports.



Note

To display the congestion statistics for the QoS profiles on a port, use the `show ports port_list qosmonitor {congestion} {no-refresh} {refresh}` command.



Note

On a V300 bridge port extender (BPE) port, neither `show ports port_list qosmonitor {congestion} {no-refresh | refresh} {port-number}` nor `show ports port_list congestion {no-refresh | refresh} {port-number}` increments when there is egress congestion. Instead, you can view egress congestion using `show ports {port_list | stack-ports stacking-port-list} rxerrors {no-refresh | refresh}` on the ingress port (upstream port in many cases).

Example

The following example shows the packets dropped due to congestion for all ports in real time:

```
# show ports congestion
Port Congestion Monitor
Port Link Packet
State Drop
=====
1:1 R 0
1:2 R 0
1:3 A 96
1:4 R 0
2:1 R 0
2:2 A 28513
2:3 R 0
2:4 R 0
2:5 R 0
2:6 R 0
```

```

2:7      R      0
2:8      R      0
3:1      R      0
3:2      R      0
3:3      R      0
3:4      R      0
=====
> indicates Port Display Name truncated past 8 characters
Link State: A-Active, R-Ready, NP-Port Not Present, L-Loopback
0->clear counters  U->page up  D->page down  ESC->exit

```

The following example shows a snapshot display of the packets dropped due to congestion for all ports:

```

# show ports congestion no-refresh
Port      Link      Packet
State     Drop
=====
1:1      R      0
1:2      R      0
1:3      A      96
1:4      R      0
2:1      R      0
2:2      A      28513
2:3      R      0
2:4      R      0
2:5      R      0
2:6      R      0
2:7      R      0
2:8      R      0
3:1      R      0
3:2      R      0
3:3      R      0
3:4      R      0
5:1      R      0
=====
> indicates Port Display Name truncated past 8 characters
Link State: A-Active, R-Ready, NP-Port Not Present, L-Loopback

```

History

This command was first available in ExtremeXOS 12.2.2.

The **refresh** keyword was added in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show ports eee

```
show ports port_list eee {no-refresh | refresh } { port-number }
```

Description

Displays the EEE statistics for the specified ports.

Syntax Description

ports	Specifies ports, or slots and ports, for which EEE information appears.
<i>port_list</i>	Optionally, specifies the list of ports, or slots and ports, for which EEE information appears.
no-refresh	Page-by-page display without continuous refresh.
refresh	Continuous refresh of output.
port-number	Displays port number instead of display string.

Default

N/A.

Usage Guidelines

This command shows various EEE statistics for the specified ports.

Example

The following command displays EEE information:

```
show ports 3,4 eee
Port EEE Statistics Monitor                               Sat Dec 29 04:10:53 2012
Port      Link  EEE      Rx      Rx      Tx      Tx
          State State    Events  Duration (uS)  Events  Duration (uS)
=====
3         A    E         1      4515403         1      4515460
4         A    E         1      4515948         1      4515966
=====

> indicates Port Display Name truncated past 8 characters
Link State: A-Active, R-Ready, NP-Port Not Present, L-Loopback
EEE State: E-Enabled, D-Disabled, NA-Not Available
0->Clear Counters  U->page up  D->page down  ESC->exit
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on all platforms that support EEE.

show ports flow-control

```
show ports port_list flowcontrol {rx-pauses} {tx-pauses} {no-refresh |
refresh}
```

Description

Displays the pause control frames received or transmitted.

Syntax Description

rx-pauses	Displays pause control frames received.
tx-pauses	Displays pause control frames transmitted.
no-refresh	Specifies a static snapshot of data.
refresh	Specifies a continuous refresh of output.

Default

By default, the pause control frames received are displayed.

Usage Guidelines

If you do not specify a port number or range of ports, the system displays information for all ports.

If you specify the no-refresh parameter, the system displays a snapshot of the data at the time you issue the command. If you do not specify the no-refresh option, the display refreshes until you press [ESC].

Example

The following example displays the flow control for ports 1, 5, and 9.

```
# show ports 1,5,9 flow-control
Flow Control Frame Monitor                               Sat Aug 18 19:35:12 2012
Port      Pause    PFC0    PFC1    PFC2    PFC3    PFC4    PFC5    PFC6    PFC7
          Rcvls    Rcvls    Rcvls    Rcvls    Rcvls    Rcvls    Rcvls    Rcvls    Rcvls
=====
1          -        -        -        - 1234567 1234567        - 1234567        -
5          -        -        -        - 1234567 1234567        - 1234567        -
16:104 1234567        -        -        -        -        -        -        -
=====
">" Name truncated, "-" rx-pause not enabled, "." Counter not available
Spacebar->Toggle screen 0->Clear counters U->Pageup D->Pagedown ESC->exit
```



Note

Use the **[spacebar]** to toggle this real-time display for all ports from received frames to transmitted frames, in that order.

History

This command was first available in ExtremeXOS 15.6.2.

The **refresh** keyword was added in ExtremeXOS 16.1.

Platform Availability

This command is supported on all platforms.

show port forward-error-correction

```
show port port_list forward-error-correction {refresh | no-refresh }
      port-number
```

Description

Shows the IEEE Forward Error Correction (FEC) Clause 74 or 91 status of the port(s).

Syntax Description

<i>port_list</i>	List of ports to show FEC modes status on.
forward-error-correction	Specifies port FEC mode status.
refresh	Specifies a continuous refresh of output.
no-refresh	Specifies a static snapshot of data.
port-number	Displays port number instead of display string.

Default

N/A.

Example

The following example shows the FEC status for port 1:

```
show port 1 forward-error-correctionPort Forward-Error-Correction Status Monitor
Port      Link    Forward-Error-Correction
          State   State      Clause
=====
1         R      NA          CL74
=====
Link State: A-Active, R-Ready, NP-Port Not Present L-Loopback
Forward-Error-Correction State: E-Enabled, D-Disabled, NA-Not Available
Clause: CL74-(Clause 74), CL91-(Clause 91), Unknown
```

History

This command was first available in ExtremeXOS 22.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show ports group

```
show ports group {port_group}
```

Description

This command displays port group information.

Syntax Description

<i>port_group</i>	Specifies the port group name.
-------------------	--------------------------------

Default

N/A.

Usage Guidelines

None.

Example

```
# show ports group fldGroupA
Group Name                Ports
-----
fldGroupA                 1-3

#show ports group
Group Name                Ports
-----
qosGroupA                 1-10
qosGroupB                 11-20
ingMeterGrpA              1-20
fldGroupA                 1-20
grp45678901234567890123456789012  3-24
```

History

This command was first available in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show port information

```
show port {mgmt | port_list | tag tag} information {detail}
```

Description

Displays detailed system-related information.

Syntax Description

mgmt	Specifies the management port.
<i>port_list</i>	Specifies one or more ports of slots and ports.
<i>tag</i>	Specifies an 802.1Q or 802.1ad tag value.
detail	Specifies detailed port information.

Default

N/A.

Usage Guidelines

This command displays information, including the following:

- Port number
- Port configuration
 - Virtual router
 - Type of port
 - LRM/MACsec Adapter port state:
 - Initializing
 - Initialization failed
 - Configuring
 - Configuration failed
 - Upgrading MCU firmware
 - Upgrading PHY firmware
 - Ready
 - Operational
 - Admin state
 - Link state and speed
 - Link Up/Down Transition
 - VLAN configuration
 - STP configuration
 - Trunking, or load sharing
 - EDP
 - ELSM (disabled; or if enabled, the ELSM link state is shown as well)
 - Load balancing
 - Learning
 - Link-flap detection status
 - Egress flooding

- Jumbo frames
- Link port up/down traps
- Port isolation status
- QoS profiles
- VMAN status
- Smart Redundancy status
- SRP status
- Additional platform-specific information
- Port partition information
- IEEE Forward Error Correction (FEC) Clause 74 or 91 modes status

If you do not specify a port number, range of ports, or tag value, detailed system-related information is displayed for all ports. The data is displayed in a table format.

The tag value may be associated with either a VMAN or a VLAN.

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

The **detail** parameter is used to provide more specific port information. The data is called out with written explanations versus displayed in a table format.



Note

The keyword **detail** displays slightly different information depending on the platform and configuration you are working with.

The link filter counter displayed with the detail keyword is calculated at the middle layer on receiving an event. The link filter up indicates the number of link transitions from down to up at the middle layer filter.

Example

The following command displays port system-related information on a switch:

```
# show ports 25 information
```

Port	Flags	Link State	ELSM/ OAM/ MACsec	Link Num UPS	Num STP	Num VLAN	Num Proto	Jumbo Size	QOS profile	Load Master
25	Em-----fMB-S-x-	active	- / -/up	3	0	1	1	9216	none	

```
> indicates Port Display Name truncated past 8 characters
Flags : a - Load Sharing Algorithm address-based,
        b - Rx and Tx Flow Control Enabled, B - Broadcast Flooding Enabled,
        D - Port Disabled, e - Extreme Discovery Protocol Enabled,
        E - Port Enabled, f - Unicast Flooding Enabled,
        F - Priority Flow Control Enabled, G - MLAG Enabled, i - Isolation,
        j - Jumbo Frame Enabled, l - Load Sharing Enabled,
        L - Extreme Link Status Monitoring Enabled,
        m - MAC Learning Enabled, M - Multicast Flooding Enabled,
        n - Ingress TOS Enabled, o - Dot1p Replacement Enabled,
        O - Ethernet OAM Enabled, p - Load Sharing Algorithm port-based,
        P - Software redundant port (Primary),
        R - Software redundant port (Redundant),
```

```
s - diffserv Replacement Enabled, S - MACsec Enabled,
v - Vman Enabled, w - MACLearning Disabled with Forwarding,
x - Rx Flow Control Enabled
```

The following command displays detailed information for port 25 on a switch:

```
# show ports 25 information detail
Port: 25
  Virtual-router: VR-Default
  Type:          UTP
  Random Early drop:      Unsupported
  Admin state:   Enabled with auto-speed sensing auto-duplex
  MACsec Link State:      Up
  Link State:    Active, 1Gbps, full-duplex
  Link Ups:      3          Last: Fri Dec 08 16:26:55 2017
  Link Downs:    1          Last: Fri Dec 08 16:18:12 2017

  VLAN cfg:
    Name: MACsec_1000, Internal Tag = 1000, MAC-limit = No-limit, Virtual
router:  VR-Default
  STP cfg:

  Protocol:
    Name: MACsec_1000 Protocol: ANY      Match all protocols.
  Trunking:      Load sharing is not enabled.

  EDP:           Disabled

  EEE:           Disabled
  ELSM:          Disabled
  Ethernet OAM:   Disabled
  MACsec:        Enabled
```

The following example displays the current status of non-Extreme optical devices:

```
# show ports 3 information detail
Port: 3
  Virtual-router: VR-Default
  Type:          Q+SR4          (Licensed)          or
                                (Unlicensed)
or
                                (Restricted)
  Random Early drop:      Unsupported
  Admin state:   Enabled
  Link State:    Ready
  Link Ups:      0          Last: --
  Link Downs:    0          Last: --
  VLAN cfg:
    Name: Default, Internal Tag = 1, MAC-limit = No-limit, Virtual router:
VR-Default
  STP cfg:
    s0(disable), Tag=(none), Mode=802.1D, State=FORWARDING
  Protocol:
    Name: Default Protocol: ANY      Match all protocols.
  Trunking:      Load sharing is not enabled.
  EDP:           Enabled
  ELSM:          Disabled
  Ethernet OAM:   Disabled
  Learning:      Enabled
  Unicast Flooding: Enabled
  Multicast Flooding: Enabled
  Broadcast Flooding: Enabled
  Jumbo:         Disabled

  Flow Control:  Rx-Pause: Enabled      Tx-Pause:      Disabled
```

```

Priority Flow Control: Disabled
Reflective Relay:      Disabled
Link up/down SNMP trap filter setting:  Enabled
Egress Port Rate:      No-limit (Restricted, 10 Gbps maximum)
Broadcast Rate:        No-limit
Multicast Rate:        No-limit
Unknown Dest Mac Rate: No-limit
QoS Profile:           None configured
Ingress Rate Shaping :  Unsupported
Ingress IPTOS Examination: Disabled

```

The following output displays show port info detail that has port-specific tag.

```

Port: 3

Virtual-router: None
Type:          BASET
Random Early drop:  Unsupported
Admin state:     Enabled with 10G full-duplex
Link State:      Active, 1Gbps, full-duplex
Link Ups:        0          Last: Wed Jun 05 00:29:24 2013
Link Downs:      0          Last: --

VLAN cfg:        Name: test, 802.1Q Tag = 200, MAC-limit = No-limit,
Virtual router:   VR-Test
    Port-specific VLAN ID: 10, 11, 200
Name: test2, 802.1Q Tag = 300, MAC-limit = No-limit,
Virtual router:   VR-Default
STP cfg:
Protocol:
Trunking:         Load sharing is not enabled.

EDP:              Disabled
ELSM:             Disabled
Ethernet OAM:      Disabled
Learning:         Enabled
Unicast Flooding:  Enabled
Multicast Flooding: Enabled
Broadcast Flooding: Enabled
Jumbo:            Enabled, MTU= 9216

Flow Control:     Rx-Pause: Enabled      Tx-Pause: Disabled
Priority Flow Control: Disabled
Reflective Relay:  Disabled
Link up/down SNMP trap filter setting:  Enabled
Egress Port Rate:  No-limit
Broadcast Rate:    No-limit
Multicast Rate:    No-limit
Unknown Dest Mac Rate: No-limit
QoS Profile:       None configured
Ingress Rate Shaping :  Unsupported
Ingress IPTOS Examination: Disabled
Ingress 802.1p Examination: Enabled
Ingress 802.1p Inner Exam: Disabled
Egress IPTOS Replacement: Disabled
Egress 802.1p Replacement: Disabled
NetLogin:         Disabled
NetLogin port mode:  Port based VLANs
Smart redundancy:   Enabled
Software redundant port: Disabled
IPFIX:            Disabled      Metering: Ingress, All Packets, All

Traffic
    IPv4 Flow Key Mask:      SIP: 255.255.255.255      DIP:
255.255.255.255
    IPv6 Flow Key Mask:      SIP:

```

```

ffff:ffff:ffff:ffff:ffff:ffff:ffff:ffff
DIP:    ffff:ffff:ffff:ffff:ffff:ffff:ffff:ffff

Far-End-Fault-Indication:    Disabled
Shared packet buffer:       default
VMAN CEP egress filtering:   Disabled
Isolation:                   Off
PTP Configured:              Disabled
Time-Stamping Mode:          None
Configured Port Partition:   4x10G
Current Port Partition:      4x10G
Available Port Partitions:   1x40G, 4x10G, 1x100G, 2x50G, 4x25G
Synchronous Ethernet:        Unsupported
Dynamic VLAN Uplink:          Disabled
VM Tracking Dynamic VLANs:    Disabled
Forward-Error-Correction:    Disabled

```

The following example shows LRM/MACsec Adapter port state information (in bold):

```

# show ports 3 information detail
Port: 3
Virtual-router:    VR-Default
Type:              SF+_CXlm (Unsupported) (LRM/MACsec Adapter)
LRM/MACsec Adapter port state:    Operational

Random Early drop:    Unsupported
Admin state:          Enabled with 10G full-duplex
Link State: Active, 10Gbps, full-duplex
Link Ups:             1                      Last: Fri Aug 16 12:07:01 2019
Link Downs:           0                      Last: --

VLAN cfg:
Name: Default, Internal Tag = 1, MAC-limit = No-limit, Virtual router:
VR-Default
STP cfg:
s0(enable), Tag=(none), Mode=802.1D, State=FORWARDING

Protocol:
Name: Default          Protocol: ANY          Match all protocols.
Trunking: Load sharing is not enabled.

EDP: Enabled

ELSM: Disabled
Ethernet OAM: Disabled
MACsec: Disabled
Learning: Enabled
Link-Flap Detection: Disabled
Unicast Flooding: Enabled
Multicast Flooding: Enabled
Broadcast Flooding: Enabled
Jumbo: Disabled
Flow Control: Rx-Pause: Enabled Tx-Pause: Disabled
Priority Flow Control: Disabled
Reflective Relay: Disabled
Link up/down SNMP trap filter setting: Enabled
Egress Port Rate: No-limit
Broadcast Rate: No-limit
Multicast Rate: No-limit
Unknown Dest Mac Rate: No-limit
QoS Profile: None configured
Ingress Rate Shaping : Unsupported
Ingress IPTOS Examination: Disabled
Ingress 802.1p Examination: Enabled

```

```

Ingress 802.1p Inner Exam:    Disabled
Ingress 802.1p Priority:      0
Egress IPTOS Replacement:    Disabled
Egress 802.1p Replacement:    Disabled
NetLogin:                    Disabled
NetLogin port mode:           Port based VLANs
Smart redundancy:             Enabled
Software redundant port:      Disabled
IPFIX:    Disabled           Metering:  Ingress, All Packets, All Traffic
          IPv4 Flow Key Mask:  SIP: 255.255.255.255          DIP: 255.255.255.255
          IPv6 Flow Key Mask:  SIP: ffff:ffff:ffff:ffff:ffff:ffff:ffff:ffff
                                DIP: ffff:ffff:ffff:ffff:ffff:ffff:ffff:ffff

Far-End-Fault-Indication:     Disabled
Shared packet buffer:         default
VMAN CEP egress filtering:    Disabled
Isolation:                    Off
PTP Configured:               Disabled
Time-Stamping Mode:           None

Dynamic VLAN Uplink:           Disabled
VM Tracking Dynamic VLANs:     Disabled

```

Example

The following command displays the indication that a 25G lane-based optic has been inserted into a port partitioned for 10G lanes (lines 4-5):

```

(pacman debug) switch-model-EXOS.6 # show port 25 information detail
Port:    25
  Virtual-router: VR-Default
  Type:          Q28+SWDM4 (Mismatches Partition)
  Random Early drop:  Unsupported
  Admin state:    Enabled with 40G full-duplex
  Link State:     Ready
  Link Ups:       0          Last: --
  Link Downs:    0          Last: --

  VLAN cfg:
    Name: Default, Internal Tag = 1, MAC-limit = No-limit, Virtual router:
VR-Default
  STP cfg:
    s0(enable), Tag=(none), Mode=802.1D, State=DISABLED

  Protocol:
    Name: Default          Protocol: ANY          Match all protocols.
  Trunking:    Load sharing is not enabled.

  EDP:          Enabled

  ELSM:         Disabled
  Ethernet OAM: Disabled
  MACsec:       Disabled
  Learning:     Enabled
  Link-Flap Detection: Disabled
  Unicast Flooding: Enabled
  Multicast Flooding: Enabled
  Broadcast Flooding: Enabled
  Jumbo:        Disabled
  Flow Control:  Rx-Pause: Enabled          Tx-Pause: Disabled
  Priority Flow Control: Disabled
  Reflective Relay: Disabled
  Link up/down SNMP trap filter setting: Enabled

```

```

Egress Port Rate:      No-limit
Broadcast Rate:        No-limit
Multicast Rate:        No-limit
Unknown Dest Mac Rate: No-limit
QoS Profile:           None configured
Ingress Rate Shaping : Unsupported
Ingress IPTOS Examination: Disabled
Ingress 802.1p Examination: Enabled
Ingress 802.1p Inner Exam: Disabled
Ingress 802.1p Priority: 0
Egress IPTOS Replacement: Disabled
Egress 802.1p Replacement: Disabled
QOS Profile Settings:  QP1  MinBw =      0%  MaxBw =    100%  MaxBuf =    100%
Weight =      1
                        QP8  MinBw =      0%  MaxBw =    100%  MaxBuf =    100%
Weight =      1

NetLogin:              Disabled
NetLogin port mode:    Port based VLANs
Smart redundancy:      Enabled
Software redundant port: Disabled
IPFIX: Disabled
IPv4 Flow Key Mask:    SIP: 255.255.255.255      DIP:
255.255.255.255
IPv6 Flow Key Mask:    SIP: ffff:ffff:ffff:ffff:ffff:ffff:ffff:ffff
DIP: ffff:ffff:ffff:ffff:ffff:ffff:ffff:ffff

Far-End-Fault-Indication: Disabled
Shared packet buffer:  default
VMAN CEP egress filtering: Disabled
Isolation:             Off
PTP Configured:        Disabled
Time-Stamping Mode:    None
Port Partition:        1x40G
Available Port Partitions: 1x40G, 4x10G, 2x50G, 4x25G
Dynamic VLAN Uplink:    Disabled
VM Tracking Dynamic VLANs: Disabled
Forward-Error-Correction: Disabled

```

Example

The following command displays the indication that a 100G optic has been inserted into an unlicensed port (lines 4-5):

```

(pacman debug) X870-32c.2 # show port 1 info detail
Port:      1
Virtual-router: VR-Default
Type:      Q28+SWDM4 (Mismatches Partition) (Unlicensed)
Random Early drop:      Unsupported
Admin state:      Enabled with 10G full-duplex
Link State:      Ready
Link Ups:      0      Last: --
Link Downs:      0      Last: --

VLAN cfg:
STP cfg:

Protocol:
Trunking:      Load sharing is not enabled.

EDP:      Enabled

ELSM:      Disabled

```

```

Ethernet OAM:      Disabled
MACsec:           Disabled
Learning:         Enabled
Link-Flap Detection: Disabled
Unicast Flooding:  Enabled
Multicast Flooding: Enabled
Broadcast Flooding: Enabled

```

Example

The following command displays port information with ingress filtering for MAC-based VLANs enabled:

```

# show ports 1:2 information detail
Port: 1:2
  Virtual-router: VR-Default
  Type: UTP
  Random Early drop: Unsupported
  Admin state: Enabled with auto-speed sensing auto-duplex
  Link State: Ready
  Link Ups: 0 Last: --
  Link Downs: 0 Last: --

  VLAN cfg:
    Name: SYS_VLAN_0300, Internal Tag = 300 (MAC-Based with ingress filtering
on), MAC-limit = No-limit, Virtual router: VR-Default
    Name: SYS_VLAN_0301, Internal Tag = 301 (MAC-Based with ingress filtering
on), MAC-limit = No-limit, Virtual router: VR-Default
    Name: SYS_VLAN_0500, 802.1Q Tag = 500, MAC-limit = No-limit, Virtual
router: VR-Default
      Port-specific VLAN ID: 500
      Name: SYS_VLAN_0501, 802.1Q Tag = 501, MAC-limit = No-limit, Virtual
router: VR-Default
      Port-specific VLAN ID: 501

  STP cfg:

  Protocol:
    Name: SYS_VLAN_0301 Protocol: ANY Match all protocols.
    Name: SYS_VLAN_0300 Protocol: ANY Match all protocols.
  Trunking: Load sharing is not enabled.

  EDP: Enabled

  EEE: Disabled
  ELSM: Disabled
  Ethernet OAM: Disabled
  MACsec: Disabled
  Learning: Enabled
  Link-Flap Detection: Disabled
  Unicast Flooding: Enabled
  Multicast Flooding: Enabled
  Broadcast Flooding: Enabled
  Jumbo: Disabled
  Flow Control: Rx-Pause: Enabled Tx-Pause: Disabled
  Priority Flow Control: Disabled
  Reflective Relay: Disabled
  Link up/down SNMP trap filter setting: Enabled
  Egress Port Rate: No-limit
  Broadcast Rate: No-limit
  Multicast Rate: No-limit
  Unknown Dest Mac Rate: No-limit
  QoS Profile: None configured
  Ingress Rate Shaping : Unsupported
  Ingress IPTOS Examination: Disabled

```

```

Ingress 802.1p Examination:      Enabled
Ingress 802.1p Inner Exam:      Disabled
Ingress 802.1p Priority:        0
Egress IPTOS Replacement:       Disabled
Egress 802.1p Replacement:      Disabled
QOS Profile Settings:           QP1  MinBw =      0%  MaxBw =    100%  MaxBuf =    100%
Weight =      1
                                QP8  MinBw =      0%  MaxBw =    100%  MaxBuf =    100%
Weight =      1

NetLogin:                       Disabled
NetLogin port mode:             MAC based VLANs with ingress filtering on
Smart redundancy:               Enabled
Software redundant port:        Disabled
IPFIX: Disabled                  Metering: Ingress, All Packets, All Traffic
                                IPv4 Flow Key Mask: SIP: 255.255.255.255      DIP:
255.255.255.255
                                IPv6 Flow Key Mask: SIP: ffff:ffff:ffff:ffff:ffff:ffff:ffff:ffff
                                DIP: ffff:ffff:ffff:ffff:ffff:ffff:ffff:ffff

auto-polarity:                  Enabled
Shared packet buffer:           default
VMAN CEP egress filtering:      Disabled
Isolation:                      Off
PTP Configured:                 Disabled
Time-Stamping Mode:             None

Dynamic VLAN Uplink:            Disabled
VM Tracking Dynamic VLANs:      Disabled

```

History

This command was first available in ExtremeXOS 10.1.

Information on ingress rate shaping was added in ExtremeXOS 11.0.

Network Login, Smart Redundancy, and rate limiting were added in ExtremeXOS 11.1.

Information on unicast, multicast, and broadcast flooding; the Port not present parameter; and autopolarity status were added in ExtremeXOS 11.2.

The netlogin parameters were added in ExtremeXOS 11.3.

The output command was modified in ExtremeXOS 12.1 so that when learning is disabled with the disabled learning port command, a new w flag appears in the output.

Link Ups and Link Downs information was added to the output and the tag value was added to the command syntax in ExtremeXOS 12.4.4.

Industrial Temperature detail added to Type field in ExtremeXOS 15.1.2.



Note

The Industrial Temperature detail is only displayed only when detail is used. Without it, the output is compressed and the optic type is not displayed.

The show output was enhanced to display the SNMP ifMib ifAlias accessible string size information.

The show output was updated to include the current status of non-Extreme optical devices in ExtremeXOS 15.4.

Link-flap detection information was added in ExtremeXOS 22_1.

Available Port Partitions was added to show port speed licensing information for ExtremeSwitching X870-96x-8c switches in ExtremeXOS 22.2.

IEEE Forward Error Correction (FEC) Clause 74 or 91 modes status was added in ExtremeXOS 22.3.

MACsec flag (S) and LRM/MACsec adapter in the detail **Type** field was added in ExtremeXOS 30.1.

LRM/MACsec Adapter port state information was added in ExtremeXOS 30.4.

Mismatched port compatibility indication was added in ExtremeXOS 31.3.

Ingress filtering for MAC-based VLANs was added in ExtremeXOS 31.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show ports link-flap-detection

```
show ports [port_list | all] link-flap-detection [disabled-ports |
configuration | counters {refresh | no-refresh}]
```

Description

Displays the link-flap detection status for ports.

Syntax Description

ports	Physical ports.
<i>port_list</i>	List of ports that you want to see link-flap detection information.
all	Shows link-flap detection information for all ports in the system.
disabled-ports	Displays ports disabled due to excessive link flapping.
configuration	Displays link-flap detection configuration information: link-flap detection enabled status, link-flap detection actions enabled, threshold value, interval value.
counters	Displays counters related to link flapping.
refresh	Continuous refresh of output (default).
no-refresh	Page by page display without continuous refresh.

Default

Output is refreshed by default.

Examples

The following example shows link-flap detection configuration for all ports:

```
show ports all link-flap-detection configuration
Maximum threshold for interval of 10 seconds : 200
Port      Link-Flap Actions  Threshold      Interval      Disable Time
Detection                                     (secs)        (secs)
=====
1          E          D--           20             6             200
2          E          D--           20             6             200
3          D          D-T           20 indefinitely until-cleared
4          E          D-T           20             5             200
5          D          D-T           20             5             200
6          D          D-T           10             5             300
...
> indicates Port Display Name truncated past 8 characters
Link-Flap Detection: (E) Enabled, (D) Disabled
Actions: (D) Disable port, (L) Generate log, (T) Generate SNMP trap
```

The following example shows link-flap detection counter information for all ports:

```
show ports all link-flap-detection counters
Port Link-Flap Detection Counters
Port      Port      Current      Total      Time      Threshold
State Link Flaps Link Flaps Elapsed Exceeded
              (secs)
=====
1          E          0           0           0           0
2          E          0           0           0           0
3          D          0           0           0           0
4          D F        21          21          4           1
> indicates Port Display Name truncated past 8 characters
Port State: E-Enabled, D-Disabled, F-Disabled by link-flap detection
0->Clear Counters U->page up D->page down ESC->exit
```

The following example shows ports disabled due to excessive link flapping:

```
show ports all link-flap-detection disabled-ports
Ports disabled due to excessive link flapping
Disabled      Remaining
Port(s)      Disable Time
              (secs)
=====
1            until-cleared
4            150
> indicates Port Display Name truncated past 15 characters
```

History

This command was first available in ExtremeXOS 22.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show ports link-scan

```
show ports link-scan {slot [slot | all]}
```

Description

Use this command to display the link scan attributes for polling port status.

Syntax Description

slot	Slot number (default all slots).
<i>slot</i>	Specifies the slot number.
all	Configures all slots.

Default

All slots.

Usage Guidelines

None.

Example

```
# show ports link-scan
Slot      Interval (ms)
-----
 1         50 (default)
 2        300
 3         50 (default)
 4         50 (default)
 5
 6
 7
 8        200
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show ports macsec-engines

```
show ports macsec-engines [qosmonitor | congestion] {no-refresh |  
refresh}
```

Description

Displays transmitted and dropped (congested) packets for each MACsec engine. Each slot has four engines: two encryption and two decryption.

Syntax Description

macsec-engines	Specifies the display of transmitted and dropped packets for each MACsec engine.
congestion	Specifies the display of packets dropped at ingress due to port congestion.
no-refresh	Specifies a static snapshot of data instead of the default dynamic display.
refresh	Specifies a continuous refresh of output.

Default

N/A.

Usage Guidelines

The MACsec engines correspond to the internal MACsec ports and not to any external switch ports.

Example

The following example displays the number of dropped packets due to congestion for each MACsec engine:

```
# show ports macsec-engines qosmonitor congestion no-refresh
MACsec Encryption/Decryption Qos Monitor
      QP1      QP2      QP3      QP4      QP5      QP6      QP7      QP8
      Pkt      Pkt      Pkt      Pkt      Pkt      Pkt      Pkt      Pkt
Engine Port  Cong  Cong  Cong  Cong  Cong  Cong  Cong  Cong
=====
1  Encryption      0      0      0      0      0      0      0      0
2                   0      0      0      0      0      0      0      0
3  Decryption      0      0      0      0      0      0      0      0
4                   0      0      0      0      0      0      0      0
```

History

This command was first available in ExtremeXOS 31.5.

Platform Availability

This command is available on ExtremeSwitching 5420 series switches.

show ports packet

```
show ports {mgmt | port_list | tag tag} packet {no-refresh | refresh}
```

Complete

Displays a snapshot or real-time histogram of packet statistics.

Syntax Description

mgmt	Specifies the management port.
<i>port_list</i>	Specifies one or more ports or slots and ports.
<i>tag</i>	Specifies an 802.1Q or 802.1ad tag value.
no-refresh	Specifies a static snapshot of data.
refresh	Specifies a continuous refresh of output.

Default

Real-time statistics.

Usage Guidelines

If you do not specify a port number, range of ports, or tag value, the system displays information for all ports; if you specify the no-refresh parameter, the system displays a snapshot of the data at the time you issue the command. To clear the counters, use the `clear counters ports` command.

The tag value may be associated with either a VMAN or a VLAN.

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 state
- Packet size

Example

The following command displays packet statistics ports 1 through 8 on a switch:

```
show ports 1-8 packet
```

Following is sample output from this command:

Port	Link	Packet Sizes					
State	0-64	65-127	128-255	256-511	512-1023	1024-1518	Jumbo
=====							
1 R	0	0	0	0	0	0	0

```

2 R      0      0      0      0      0      0      0
3 R      0      0      0      0      0      0      0
4 R      0      0      0      0      0      0      0
5 R      0      0      0      0      0      0      0
6 R      0      0      0      0      0      0      0
7 R      0      0      0      0      0      0      0
8 R      0      0      0      0      0      0      0
=====
Link State: A-Active, R-Ready, NP-Port not present, L-Loopback

```

History

This command was first available in ExtremeXOS 10.1.

The Port Not Present variable was added in ExtremeXOS 11.2.

The **no-refresh** variable was added in ExtremeXOS 11.3.

The tag value was added in ExtremeXOS 12.4.4.

The **refresh** keyword was added in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show ports partition-template

```
show ports partition-template {slot [ slot | all ] }
```

Description

This command shows the bandwidth usage of the six uplink ports (49, 53, 57, 61, 65, 69) of the ExtremeSwitching X690 series switches.

Syntax Description

partition-template	Specifies showing the bandwidth usage of the ExtremeSwitching X690 series switches' uplink ports.
slot	Specifies a slot number.
<i>slot</i>	Designates the slot number.
all	Shows partition for all of the six uplink ports (default).

Default

If you do not designate slots, all slots are shown.

Usage Guidelines

ExtremeSwitching X690 series switches allow 400Gbps of combined I/O bandwidth for the six uplink ports (49, 53, 57, 61, 65, 69). If all of these uplink ports are configured to use their maximum capacity (49, 53 can operate at 40G; and 57, 61, 65, and 69 can operate at 100G), they can exceed the total allowed bandwidth.

The default port partition template configures the six ports as two 100G and four 40G ports. This default uses all six ports within the available total I/O bandwidth. You can also create a configuration of four 100G uplink ports. In this configuration, two QSFP ports, 49 and 53, are unused so as not to exceed the total allowed I/O bandwidth. Unused ports appears as "NP" (not present).

Example

The following example shows the bandwidth usage of the six uplink ports on an ExtremeSwitching X690 switch:

```
show ports partition-template all
Slot   Current Partition Template
-----
1      4x100G
      49-52:  NP, 53-56:  NP, 57-60: 100G, 61-64: 100G, 65-68: 100G, 69-72: 100G

Slot   Configured Partition Template
-----
1      4x100G
      49-52:  NP, 53-56:  NP, 57-60: 100G, 61-64: 100G, 65-68: 100G, 69-72: 100G
```

History

This command was first available in ExtremeXOS 22.3.

Platform Availability

This command is available on the ExtremeSwitching X690 series switches.

show ports protocol filter

```
show ports [port_list | all] protocol filter {detail}
```

Description

Displays the protocol filtering configuration and status.

Syntax Description

<i>port_list</i>	Displays port list, separated by a comma (,)or dash (-).
all	Displays all ports.
detail	Displays detailed configuration and status.

Default

Displays all protocol filters.

Usage Guidelines

Use this command to display the protocol filtering configuration and status.

Example

The following example displays the filtering configuration and status for ports 1-4:

```
# show ports 1-4 protocol filter
Port Protocol   Destination      Protocol Id  Field  Field Field  Packets
#   Filter Name Address          Type  Value Offset Value  Mask  Filtered
-----
1  my_list      01:80:C2:00:00:02 etype 0x8902   14   03:04  FF:FF
2300
      01:80:C2:00:00:00 snap 0x4041                5000
      01:80:C2:00:00:00 snap 0x4041   16   01:02> FF:FF> 5000
2  lacp         01:80:C2:00:00:02 etype 0x8902   14    01    FF    3200
3  (none)
4  (none)

> indicates that the value was truncated to the column size in the output. Use the
"detail" option to see the complete value.
```

The following example displays output for the `show ports protocol filter detail` command:

```
show ports 1-4 protocol filter detail
Port 1
  Protocol Filter Name: my_list
  Destination Address : 01:80:C2:00:00:02
  Protocol Id Type    : etype
  Protocol Id Value    : 0x8902
  Field Offset        : 14
  Field Value         : 03:04
  Field Mask          : FF:FF
  Packets Filtered    : 2300
  Destination Address : 01:80:C2:00:00:00
  Protocol Id Type    : snap
  Protocol Id Value    : 0x4041
  Field Offset        : 16
  Field Value         : 01:02:03:04
  Field Mask          : FF:FF:FF:FF
  Packets Filtered    : 5000
  Destination Address : 01:80:C2:00:00:00
  Protocol Id Type    : snap
  Protocol Id Value    : 0x4041
  Field Offset        :
  Field Value         :
  Field Mask          :
  Packets Filtered    : 5000

Port 2
  Protocol Filter Name: lacp
  Destination Address : 01:80:C2:00:00:02
  Protocol Id Type    : etype
  Protocol Id Value    : 0x8902
  Field Offset        : 14
```



```

Field Value      : 01
Field Mask       : FF
Packets Filtered : 3200

Port 3
  Protocol Filter Name: (none)

Port 4  Protocol Filter Name: (none)

```

History

This command was first available in ExtremeXOS 15.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show ports qosmonitor

```

show ports port_list qosmonitor {ingress | egress} {bytes | packets}
        {no-refresh | refresh} {port-number}

```

Description

Displays egress traffic counts or ingress traffic counts for each *QoS* profile on the specified ports.

Syntax Description

<i>port_list</i>	Specifies one or more slots and ports.
egress	Specifies the display of egress traffic counts. Default.
bytes	Specifies to display ingress or egress traffic counts in bytes.
packets	Specifies to display ingress or egress traffic counts in packets.
no-refresh	Specifies a static snapshot of data instead of the default dynamic display.
refresh	Specifies a continuous refresh of output.
port-number	Displays port number instead of display string.

Default

Displays egress packet counts in real-time.

Usage Guidelines

The bottom line in the real-time display shows keys that you can press to change the display. For example, the spacebar toggles the display between QoS traffic counts in either packets or bytes.

If you specify the no-refresh parameter, the system displays a snapshot of the data at the time you issue the command.

If you do not specify a port number or range of ports when displaying ingress or egress traffic counts, traffic counts are displayed for all ports.

Example

The following example shows the egress packet counts for the specified ports:

```
# show ports 1:1-1:2 qosmonitor
Qos Monitor Req Summary
Port      QP1      QP2      QP3      QP4      QP5      QP6      QP7      QP8
Pkt       Pkt       Pkt       Pkt       Pkt       Pkt       Pkt       Pkt
Xmts       Xmts       Xmts       Xmts       Xmts       Xmts       Xmts       Xmts
=====
1:1         0         0         0         0         0         0         0         0
1:2         0         0         0         0         0         0         0         0
=====
> indicates Port Display Name truncated past 8 characters
Spacebar->Toggle screen 0->Clear counters  U->Page up  D->Page down ESC->exit
```

The following example shows the ingress packet counts for the specified ports:

```
# show ports 1:1-1:2 qosmonitor ingress
Qos Monitor Req Summary
Port      IQP1     IQP2     IQP3     IQP4     IQP5     IQP6     IQP7     IQP8
Pkt       Pkt       Pkt       Pkt       Pkt       Pkt       Pkt       Pkt
Xmts       Xmts       Xmts       Xmts       Xmts       Xmts       Xmts       Xmts
=====
1:1         0         0         0         0         0         0         0         0
1:2         0         0         0         0         0         0         0         0
=====
> indicates Port Display Name truncated past 8 characters
Spacebar->Toggle screen 0->Clear counters  U->Page up  D->Page down ESC->exit
```

History

This command was first available in ExtremeXOS 10.1.

The ingress information was added in ExtremeXOS 11.0.

Also, you must specify the ports in ExtremeXOS 11.0.

The **egress** and **no-refresh** keywords were added in ExtremeXOS 11.3.

The **bytes** and **packets** keywords, as well as the toggling functionality, were added in ExtremeXOS 11.4.

The **refresh** keyword was added in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show ports qosmonitor {congestion}

```
show ports port_list qosmonitor {congestion} {no-refresh | refresh}
           {port-number}
```

Description

Displays egress packet counts or dropped-traffic counts for each [QoS](#) profile on the specified ports.

Syntax Description

<i>port_list</i>	Specifies one or more slots and ports.
congestion	Specifies the display of packets dropped at ingress due to port congestion.
no-refresh	Specifies a static snapshot of data instead of the default dynamic display.
refresh	Specifies a continuous refresh of output.
port-number	Displays port number instead of display string.

Default

Displays egress packet counts in real-time.

Usage Guidelines

The bottom line in the real-time display shows keys that you can press to change the display. For example, the spacebar toggles the display between egress packet counts and ingress dropped-packet counts.



Note

This command does not work properly if another CLI session is displaying congestion statistics in real time.

If you specify the no-refresh parameter, the system displays a snapshot of the data at the time you issue the command.

The dropped packet display is limited to the 8 most-significant digits.

When you display the packet counts for a port, this action configures the hardware to monitor that port. If the switch was previously configured to monitor a different port, the counters are reset for the new port. If the selected port is the last port displayed on the switch, the counters are not reset.



Note

On a V300 bridge port extender (BPE) port, neither `show ports port_list qosmonitor {congestion} {no-refresh | refresh} {port-number}` nor `show ports port_list congestion {no-refresh | refresh} {port-number}` increments when there is egress congestion. Instead, you can view egress congestion using `show ports {port_list | stack-ports stacking-port-list} rxerrors {no-refresh | refresh}` on the ingress port (upstream port in many cases).

Example

The following example shows the egress packet counts for the specified ports:

```
# show ports 2:1, 3:6 qosmonitor
QoS Monitor Req Summary
Port      QP1      QP2      QP3      QP4      QP5      QP6      QP7      QP8
Pkt       Pkt       Pkt       Pkt       Pkt       Pkt       Pkt       Pkt
Xmts      Xmts      Xmts      Xmts      Xmts      Xmts      Xmts      Xmts
=====
2:1        0         0         0         0         0         0         0         0
3:6        0         0         0         0         0         0         0         0
=====
> indicates Port Display Name truncated past 8 characters
Spacebar->Toggle screen 0->Clear counters  U->Page up  D->Page down  ESC->exit
```

The next example shows the dropped packet counts for the specified ports:

```
# show ports 2:1, 3:6 qosmonitor congestion
QoS Monitor Req Summary
Port      QP1      QP2      QP3      QP4      QP5      QP6      QP7      QP8
Pkt       Pkt       Pkt       Pkt       Pkt       Pkt       Pkt       Pkt
Cong      Cong      Cong      Cong      Cong      Cong      Cong      Cong
=====
2:1        0         0         0         0         0         0         0         0
3:6      8745         0        129         0         0         0         0         0
=====
> indicates Port Display Name truncated past 8 characters
Spacebar->Toggle screen 0->Clear counters  U->Page up  D->Page down  ESC->exit
```

History

This command was first available in ExtremeXOS 10.1.

You must specify the ports in ExtremeXOS 11.0.

The no-refresh keyword was added in ExtremeXOS 11.3.

The congestion keyword was added in ExtremeXOS 12.2, and the toggling functionality was modified to switch between egress packets and dropped packets.

The **refresh** keyword was added in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show ports rate-limit flood

```
show ports { port_list | port_group | tag tag} rate-limit flood {out-  
actions | out-of-profile {disabled-ports}} {no-refresh | refresh}
```

Description

Displays rate-limit discard statistics.

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
<i>port_group</i>	Port group name.
tag	Display ports in VMAN/VLANs with this IEEE 802.1Q or 802.1ad tag
<i>tag</i>	Tag ID between 1 and 4095.
rate-limit	Displays rate-limit discard statistics.
flood	Flood traffic.
out-actions	Display out-of-profile actions.
out-of-profile	Display rate-limit out-of-profile status.
disabled-ports	Display ports that have been disabled due to out-of-profile status.
no-refresh	Specifies a static snapshot of data.
refresh	Specifies a continuous refresh of output.

Default

N/A.

Usage Guidelines

This command displays the per port ingress rate-limit flood traffic counter as well as information about received packets that have not been discarded due to rate-limiting.

It is used to show the results of the `configure ports port_list rate-limit flood [broadcast | multicast | unknown-destmac] [no-limit | pps]` command.



Note

As part of the system health check, the system polls the Rate-limit Flood Counters every 5 minutes and looks for non-zero counters on a port. A HAL.RateLimit.Info log message is logged when this is first detected on a port to alert the user that something in the network has triggered the rate limiting to occur. The message is not be logged again unless the counters are cleared.

Example

Following is sample out put from this command:

```
* switch # show ports rate-limit flood
Port Rate-Limit Discard Monitor
Port Link Rx Pkt Rx Byte Rx Pkt Rx Pkt Flood Rate
State Count Count Bcast Mcast Exceeded
=====
1 A 0 0 0 0 0
2 A 0 0 0 0 0
3 R 0 0 0 0 0
4 R 0 0 0 0 0
5 R 0 0 0 0 0
6 R 0 0 0 0 0
7 R 0 0 0 0 0
8 R 0 0 0 0 0
9 R 0 0 0 0 0
10 R 0 0 0 0 0
11 R 0 0 0 0 0
12 R 0 0 0 0 0
13 R 0 0 0 0 0
14 R 0 0 0 0 0
15 R 0 0 0 0 0
16 R 0 0 0 0 0
=====
> indicates Port Display Name truncated past 8 characters
Link State: A-Active, R-Ready, NP-Port Not Present, L-Loopback
0->Clear Counters U->page up D->page down ESC->exit
```

The following command displays information without a screen refresh on a switch:

```
show ports rate-limit flood no-refresh
```

Following is sample out put from this command:

```
* switch # show ports rate-limit flood no-refresh
Port Rate-Limit Discard Monitor
Port Link Rx Pkt Rx Byte Rx Pkt Rx Pkt Flood Rate
State Count Count Bcast Mcast Exceeded
=====
1 A 0 0 0 0 0
2 A 0 0 0 0 0
3 R 0 0 0 0 0
4 R 0 0 0 0 0
5 R 0 0 0 0 0
6 R 0 0 0 0 0
```

```

7      R      0      0      0      0      0
8      R      0      0      0      0      0
9      R      0      0      0      0      0
10     R      0      0      0      0      0
11     R      0      0      0      0      0
12     R      0      0      0      0      0
13     R      0      0      0      0      0
14     R      0      0      0      0      0
15     R      0      0      0      0      0
16     R      0      0      0      0      0
17     R      0      0      0      0      0
18     R      0      0      0      0      0
19     A      0      0      0      0      0
20     A      0      0      0      0      0
21     R      0      0      0      0      0
22     R      0      0      0      0      0
23     R      0      0      0      0      0
24     R      0      0      0      0      0
25     R      0      0      0      0      0
26     R      0      0      0      0      0
27     R      0      0      0      0      0
28     R      0      0      0      0      0
=====
> indicates Port Display Name truncated past 8 characters
Link State: A-Active, R-Ready, NP-Port Not Present L-Loopback

```

The following examples use the **out-actions** and **out-of-profile** keywords:

```

# show ports rate-limit flood out-actions

Port      Flood Type      Out-actions
Log      Trap      Disable
Port
-----
1      Unknown Dest MAC      On      On      Off
1      Broadcast      On      On      Off
1      Multicast      On      On      Off
2      Unknown Dest MAC      On      Off      On
2      Broadcast      On      Off      On
2      Multicast      On      Off      On
. . .

# show ports rate-limit flood out-of-profile

Port      Flood Type      Status      Counter
-----
1      Unknown Dest MAC      Out of profile      1234567890
1      Broadcast      Ok      0
1      Multicast      Ok      0
2      Unknown Dest MAC      Ok      0
2      Broadcast      Out of profile      12345
. . .

# show ports rate-limit flood out-of-profile disabled-ports
Disabled
Port      Flood Type      Status      Counter
-----
2      Broadcast      Out of profile      12345
. . .

```

History

This command was first available in ExtremeXOS 12.2.

The *port-group*, **out-actions**, **out-of-profile**, **disabled-ports**, and **refresh** options were added in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show ports redundant

```
show ports redundant
```

Description

Displays detailed information about redundant ports.

Syntax

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays information on software-controlled redundant ports on the switch:

```
show ports redundant
```

Following is sample output from this command:

```
Primary: *1:1          Redundant: 3:1, Link on/off option: OFF  
Flags: (*)Active, (!) Disabled, (g) Load Share Group
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show ports rxerrors

```
show ports {port_list | stack-ports stacking-port-list} rxerrors {no-  
refresh | refresh}
```

Description

Displays real-time receive error statistics. The switch automatically refreshes the output unless otherwise specified.

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
<i>stacking-port-list</i>	Specifies one or more stacking ports or slots. Applies to SummitStack and the ExtremeSwitching series switches only.
no-refresh	Specifies that auto-refresh is disabled. The output provides a real-time snapshot of the receive errors at the time the command is issued. This setting is not saved.
refresh	Specifies a continuous refresh of output.

Default

The switch automatically refreshes the output.

Usage Guidelines

If you do not specify a port number or range of ports, receive error statistics are displayed for all ports.

If you do not specify the no-refresh parameter, the switch automatically refreshes the output (this is the default behavior).

If you specify the no-refresh parameter, the output provides a snapshot of the real-time receive error statistics at the time you issue the command and displays the output in page-by-page mode (this was the default behavior in ExtremeXOS 11.2 and earlier). This setting is not saved; therefore, you must specify the no-refresh parameter each time you want a snapshot of the port receive errors.

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

Collected Port Receive Error Information

The switch collects the following port receive error information:

- Port Number.
- Link State—The current state of the link. Options are:
 - Active (A)—The link is present at this port.
 - Ready (R)—The port is ready to accept a link.
 - Loopback (L)—The port is in Loopback mode.

- Receive Bad CRC Frames (RX CRC)—The total number of frames received by the port that were of the correct length, but contained a bad FCS value.
- Receive Oversize Frames (RX Over)—The total number of good frames received by the port greater than the supported maximum length of 1,522 bytes.
- Receive Undersize Frames (RX Under)—The total number of frames received by the port that were less than 64 bytes long.
- Receive Fragmented Frames (RX Frag)—The total number of frames received by the port were of incorrect length and contained a bad FCS value.
- Receive Jabber Frames (RX 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 packets dropped due to the memory queue being full.

Port Monitoring Display Keys

For information about the available port monitoring display keys, see the [show ports statistics](#) command.



Note

On a V300 bridge port extender (BPE) port, neither `show ports port_list qosmonitor {congestion} {no-refresh | refresh} {port-number}` nor `show ports port_list congestion {no-refresh | refresh} {port-number}` increments when there is egress congestion. Instead, you can view egress congestion using this command on the ingress port (upstream port in many cases).

Example

The following command displays receive error statistics for all ports on the switches with auto-refresh enabled (the default behavior):

```
show ports rxerrors
```

The following is sample truncated output from this command:

```
Port Rx Error Monitor Tue Jul  5 15:07:13 UTC 2005
Port  Link      Rx      Rx      Rx      Rx      Rx      Rx      Rx
State Crc      Over    Under   Frag    Jabber  Align   Lost
=====
1      R          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
9      R          0        0         0         0         0         0
10     R          0        0         0         0         0         0
```

```

11      R      0      0      0      0      0      0      0
12      R      0      0      0      0      0      0      0
13      R      0      0      0      0      0      0      0
14      R      0      0      0      0      0      0      0
15      R      0      0      0      0      0      0      0
16      R      0      0      0      0      0      0      0
17      R      0      0      0      0      0      0      0
=====
Link State: A-Active, R-Ready, NP-Port Not Present, L-Loopback
0->Clear Counters  U->page up  D->page down  ESC->exit

```

History

This command was first available in ExtremeXOS 10.1.

The link state information was updated to include NP-Port not present in ExtremeXOS 11.2.

Support for the auto-refresh functionality and the no-refresh parameter were added in ExtremeXOS 11.3. Auto-refresh continually updates the display. The no-refresh parameter takes a real-time snapshot of the display at the time you issue the command.

The **refresh** keyword was added in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show ports sharing

```
show ports port_list sharing distribution configuration
```

Description

Displays port load-sharing groups, or LAGs.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Output from this command displays the following information:

- Config Master—The port that is configured as the master logical port of the link aggregation group (LAG). This number is also the LAG group ID.
- Current Master—In LACP, this is the port that is currently the LAG group ID, or master logical port for the LAG.

- **Agg Control**—This is the aggregation control for the specified LAG; it can be either static, LACP or health-check. In LACP, it is the aggregation control for the specified LAG.
- **Min Active**—The minimum number of active links that must be up for the trunk to remain up.
- **Ld Share Algorithm**—The algorithm used for the link aggregation. The available link aggregation algorithms vary among platforms; see the [Switch Engine 32.3 User Guide](#) for more information.
- **Ld Share Group**—The specific ports that belong to each LAG, or the port numbers in the trunk. A port can belong to only one LAG, either static or dynamic.
- **Agg Mbr**—In LACP, this shows whether the port has been added to the aggregator or not; it will be either Y for yes or - for no.
- **Link State**—This is the current status of the link.
- **Link Up transitions**—Number of times the link has cycled through being up, then down, then up.
- **LAG distribution mode**—Displays the configured distribution mode and distribution port lists for LAGs.

Example

Following is sample output displaying link aggregation on a switch:

```
show port sharing
Load Sharing Monitor
Config      Current Agg      Min      Ld Share  Ld Share  Agg  Link  Link Up
Master      Master  Control  Active  Algorithm  Group   Mbr   State Transitions
=====
      14                                <2      L2       14      -      R      0
                                L2       15      -      A      1
                                L2       16      -      R      0
      17      35      Static    1      L2       17      -      R      0
                                L2       35      Y      A      1
=====
Link State: A-Active, D-Disabled, R-Ready, NP-Port not present, L-Loopback
Minimum Active: (<) Group is down. # active links less than configured minimum
Load Sharing Algorithm: (L2) Layer 2 address based, (L3) Layer 3 address based
                        (L3_L4) Layer 3 address and Layer 4 port based
                        (custom) User-selected address-based configuration

Custom Algorithm Configuration: ipv4 L3-and-L4, xor
Custom Hash Seed: Switch MAC address (0x969BF582)
Number of load sharing trunks: 2
```

History

This command was first available in ExtremeXOS 10.1.

The LACP feature was added in ExtremeXOS 11.3.

The Health Check LAG was added in ExtremeXOS 12.1.3.

The round-robin algorithm was added in ExtremeXOS 12.5.

Support for the Min Active parameter was added in ExtremeXOS 15.7.1.

Custom hash seed information was added in ExtremeXOS 30.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show ports stack-ports congestion

```
show ports stack-ports stack_port_list congestion {no-refresh | refresh}  
           port-number
```

Description

Displays the stacking port congestion statistics (dropped packets) for the specified ports on the front panel.

Syntax Description

<i>stack_port_list</i>	Specifies one or more stacking ports.
no-refresh	Specifies a static snapshot of data instead of the default dynamic display.
refresh	Specifies a continuous refresh of output.
port-number	Displays port number instead of display string in the output of the command.

Default

Displays the port congestion statistics for all ports in real-time.

Usage Guidelines

The bottom line in the real-time display shows keys that you can press to change the display. For example, you can clear the counters or page up or down through the list of ports.



Note

If you are displaying congestion statistics in real time and another CLI session resets the counters for a port you are monitoring, the counters displayed in your session for that port are also reset.

If you specify the no-refresh parameter, the system displays a snapshot of the data at the time you issue the command.



Note

Packets can be dropped at multiple locations along the path through the hardware. The per-port congestion counters count all dropped packets for all ports.

If you do not specify a port number or range in the command, dropped packet counts are displayed for all ports.



Note

To display the congestion statistics for the [QoS](#) profiles on a stack port, use the `show ports stack-ports stack_port_list qosmonitor congestion {no-refresh | refresh} {port-number}` command.

Example

The following example shows the packets dropped due to congestion for all ports in real time:

```
# show ports stack-ports congestion
Port Congestion Monitor                                     Fri Apr 21 13:02:37 2017
Port      Link      Packet
          State     Drop
=====
1:1       R         0
1:2       R         0
1:3       A        96
1:4       R         0
2:1       R         0
2:2       A      28513
2:3       R         0
2:4       R         0
2:5       R         0
2:6       R         0
2:7       R         0
2:8       R         0
3:1       R         0
3:2       R         0
3:3       R         0
3:4       R         0
=====
> indicates Port Display Name truncated past 8 characters
Link State: A-Active, R-Ready, NP-Port Not Present, L-Loopback
0->clear counters  U->page up  D->page down  ESC->exit
```

History

This command was first available in ExtremeXOS 22.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show ports stack-ports qosmonitor

```
show ports stack-ports stack_port_list qosmonitor {no-refresh | refresh}
{port-number}
```

Description

Displays traffic counts for each [QoS](#) profile on the specified stacking ports.

Syntax Description

<i>stack_port_list</i>	Specifies one or more stacking ports.
no-refresh	Specifies a static snapshot of data instead of the default dynamic display.
refresh	Specifies a continuous refresh of output.
port-number	Displays port number instead of display string in the output of the command.

Default

Displays packet counts in real-time.

If you do not specify a port number or range of ports, traffic counts appear for all ports.

Usage Guidelines

The bottom line in the real-time display shows keys that you can press to change the display. For example, the space bar changes the display from QoS traffic counts in either packets or bytes.

If you specify the no-refresh option, you see a snapshot of the data at the time you issued the command.

Example

The following example shows the packet counts for ports 1:1 and 1:2:

```
# show ports stack-ports 1:1-1:2 qosmonitor
Qos Monitor Req Summary                               Fri Apr 21 10:58:23 2017
Port      QP1      QP2      QP3      QP4      QP5      QP6      QP7      QP8
Pkt       Pkt       Pkt       Pkt       Pkt       Pkt       Pkt       Pkt
Xmts      Xmts      Xmts      Xmts      Xmts      Xmts      Xmts      Xmts
=====
1:1         0         0         0         0         0         0         0         0
1:2         0         0         0         0         0         0         0         0
=====
> indicates Port Display Name truncated past 8 characters
Spacebar->Toggle screen 0->Clear counters  U->Page up  D->Page down ESC->exit
```

History

This command was first available in ExtremeXOS 22.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show ports stack-ports qosmonitor congestion

```
show ports stack-ports stack_port_list qosmonitor congestion {no-refresh
| refresh} {port-number}
```

Description

Displays packet counts or dropped-traffic counts for each [QoS](#) profile on the specified stacking ports.

Syntax Description

<i>stack_port_list</i>	Specifies one or more stacking ports.
congestion	Specifies the display of packets dropped due to port congestion.
no-refresh	Specifies a static snapshot of data instead of the default dynamic display.
refresh	Specifies a continuous refresh of output.
port-number	Displays port number instead of display string in the output of the command.

Default

Displays packet counts in real-time.

If you do not specify a port number or range of ports, traffic counts appear for all ports.

Usage Guidelines

The bottom line in the real-time display shows keys that you can press to change the display. For example, the space bar changes the display from packet counts and dropped-packet counts.



Note

This command does not work properly if another CLI session is displaying congestion statistics in real time.

If you specify the no-refresh option, you see a snapshot of the data at the time you issued the command.

The dropped packet display is limited to the 8 most-significant digits.

When you display the packet counts for a port, this action configures the hardware to monitor that port. If the switch was previously configured to monitor a different port, the counters are reset for the new port. If the selected port is the last port displayed on the switch, the counters are not reset.

Example

The following example shows the dropped packet counts for ports 2:1 and 3:6:

```
# show ports 2:1, 3:6 qosmonitor congestion
QoS Monitor Req Summary                               Fri Apr 21 01:17:14 2017
Port   QP1      QP2      QP3      QP4      QP5      QP6      QP7      QP8
Pkt    Pkt    Pkt    Pkt    Pkt    Pkt    Pkt    Pkt
Cong   Cong   Cong   Cong   Cong   Cong   Cong   Cong
=====
2:1    0       0       0       0       0       0       0       0
3:6    8745    0       129     0       0       0       0       0
=====
```



```
> indicates Port Display Name truncated past 8 characters
Spacebar->Toggle screen 0->Clear counters U->Page up D->Page down ESC->exit
```

History

This command was first available in ExtremeXOS 22.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show ports statistics

```
show ports {port_list | stack-ports stacking-port-list} statistics {no-  
refresh | refresh}
```

Description

Displays real-time port statistic information. The switch automatically refreshes the output unless otherwise specified.

Syntax Description

<i>stacking-port-list</i>	Specifies one or more stacking slots and ports. Applies to SummitStack and the ExtremeSwitching series switches only.
<i>port_list</i>	Specifies one or more ports or slots and ports.
no-refresh	Specifies that auto-refresh is disabled. The output provides a real-time snapshot of the port statistics at the time the command is issued. This setting is not saved.
refresh	Specifies a continuous refresh of output.

Default

The switch automatically refreshes the output.

Usage Guidelines

If you do not specify a port number or range of ports, statistics are displayed for all ports.

If you do not specify the no-refresh parameter, the switch automatically refreshes the output (this is the default behavior).

If you specify the no-refresh parameter, the output provides a snapshot of the real-time port statistics at the time you issue the command and displays the output in page-by-page mode (this was the default behavior in ExtremeXOS 11.2 and earlier). This setting is not saved; therefore, you must specify the no-refresh parameter each time you want a snapshot of the port statistics.

Jumbo frame statistics are displayed for 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.

Collected Port Statistics

The switch collects the following port statistic information:

- Port Number.
- Link State—The current state of the link. Options are:
 - Active (A)—The link is present at this port.
 - Ready (R)—The port is ready to accept a link.
 - Loopback (L)—The port is in Loopback mode.
- 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.



Note

On switches, when a broadcast jumbo frame is sent, the RX Bcast counter is not updated. The RX Pkt counter is updated to reflect the received broadcast jumbo frames.

- Received Multicast (RX Mcast)—The total number of frames received by the port that are addressed to a multicast address.

Port Monitoring Display Keys

The following table describes the keys used to control the display that appears if auto-refresh is enabled (the default behavior).

Table 41: Port Monitoring Display Keys with Auto-Refresh Enabled

Key(s)	Description
U	Displays the previous page of ports.
D	Displays the next page of ports.
[Esc]	Exits from the screen.
O	Clears all counters.

The following table describes the keys used to control the display that appears if you auto-refresh is disabled.

Table 42: Port Monitoring Displays Keys with Auto-Refresh Disabled

Key	Description
Q	Exits from the screen.
[Space]	Displays the next page of ports.

Example

The following command displays port statistics for all ports on switches with auto-refresh enabled (the default behavior):

```
show ports statistics
```

The following is truncated sample output from this command:

```

Port Statistics
Port      Link      Tx Pkt      Tx Byte      Rx Pkt      Rx Byte      Rx      Rx
State     Count     Count      Count      Count      Count      Bcast   Mcast
=====
1         R         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
9         R         0           0           0           0           0       0
10        R         0           0           0           0           0       0
11        R         0           0           0           0           0       0
12        R         0           0           0           0           0       0
13        R         0           0           0           0           0       0
14        R         0           0           0           0           0       0
15        R         0           0           0           0           0       0
16        R         0           0           0           0           0       0
17        R         0           0           0           0           0       0
=====
Link State: A-Active, R-Ready, NP-Port Not Present, L-Loopback
0->Clear Counters  U->page up  D->page down  ESC->exitPort Statistics

```

History

This command was first available in ExtremeXOS 10.1.

The link state information was updated to include NP-Port not present in ExtremeXOS 11.2.

Support for the auto-refresh functionality and the no-refresh parameter were added in ExtremeXOS 11.3. Auto-refresh continually updates the display. The no-refresh parameter takes a real-time snapshot of the display at the time you issue the command.

The **refresh** keyword was added in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show ports transceiver information detail

```
show ports {port_list | tag tag} transceiver information detail
```

Description

Displays detailed information about the optical transceiver.

Syntax Description

<i>port_list</i>	Specifies the port number(s).
<i>tag</i>	Specifies an 802.1Q or 802.1ad tag value.

Default

N/A.

Usage Guidelines

Digital Diagnostic Monitoring Interface (DDMI) provides critical system information about the installed optical modules. Use this command to monitor the condition of XFP, SFP, and SFP+ optical transceiver modules.

The tag value may be associated with either a VMAN or a VLAN.

If you try to execute this command on one of the ports in the port list that is non-compliant with DDMI, the following error message is displayed and the command does not go through:

```
Port 3:1 This command is not supported on this port. All ports and
transceiver of the ports requested in the command need to support DDMI.
```

If you try to execute this command on one of the ports in the port list on which the transceiver is non-compliant with DDMI, the following error message is displayed:

```
Port 3:1 This media/transceiver does not support enhanced digital
diagnostic monitoring interface (DDMI). All ports and transceiver of the
ports requested in the command need to support DDMI.
```

Example

```
# show port 2:* conf
Port Configuration Monitor
Port      Virtual   Port Link  Auto   Speed      Duplex  Flow  Load  Media
          router    State State Neg   Cfg Actual Cfg Actual Cntrl Master Pri Red
=====
2:1      VR-Default E      A    OFF 40000 40000 FULL FULL  SYM      %Q_UNKWN
2:2      VR-Default E      NP   OFF 10000      FULL
2:3      VR-Default E      NP   OFF 10000      FULL
                                     NONE
                                     NONE
```

```

2:4      VR-Default  E      NP      OFF 10000      FULL      NONE
2:5      VR-Default  E      R       OFF 40000      FULL      NONE
2:6      VR-Default  E      NP      OFF 10000      FULL      NONE
2:7      VR-Default  E      NP      OFF 10000      FULL      NONE
2:8      VR-Default  E      NP      OFF 10000      FULL      NONE
2:9      VR-Default  E      R       OFF 40000      FULL      %Q+LR4
2:10     VR-Default  E      NP      OFF 10000      FULL      NONE
2:11     VR-Default  E      NP      OFF 10000      FULL      NONE
2:12     VR-Default  E      NP      OFF 10000      FULL      NONE
2:13     VR-Default  E      R       OFF 40000      FULL      %Q+SR4
2:14     VR-Default  E      NP      OFF 10000      FULL      NONE

```

```

=====
> indicates Port Display Name truncated past 8 characters
Link State: A-Active, R-Ready, NP-Port Not Present, L-Loopback
Port State: D-Disabled, E-Enabled
Media: !/$/%/* - use "show port info detail" for more information
0->Clear Counters  U->page up  D->page down ESC->exit

```



Note

When a license is not installed, the restricted transceivers use a '\$'.

```

# show port 2:* conf no
Port Configuration
Port      Virtual      Port  Link  Auto   Speed      Duplex      Flow  Load  Media
router    State State Neg   Cfg Actual   Cfg Actual Cntrl Master Pri Red
=====
2:1      VR-Default  E      R      OFF 40000      FULL      $Q+SR4
2:2      VR-Default  E      NP     OFF 10000      FULL      NONE
2:3      VR-Default  E      NP     OFF 10000      FULL      NONE
2:4      VR-Default  E      NP     OFF 10000      FULL      NONE
2:5      VR-Default  E      R      OFF 40000      FULL      NONE
2:6      VR-Default  E      NP     OFF 10000      FULL      NONE
2:7      VR-Default  E      NP     OFF 10000      FULL      NONE
2:8      VR-Default  E      NP     OFF 10000      FULL      NONE
2:9      VR-Default  E      R      OFF 40000      FULL      $Q+LR4
2:10     VR-Default  E      NP     OFF 10000      FULL      NONE
2:11     VR-Default  E      NP     OFF 10000      FULL      NONE
2:12     VR-Default  E      NP     OFF 10000      FULL      NONE
2:13     VR-Default  E      R      OFF 40000      FULL      $Q+SR4
2:14     VR-Default  E      NP     OFF 10000      FULL      NONE
2:15     VR-Default  E      NP     OFF 10000      FULL      NONE
2:16     VR-Default  E      NP     OFF 10000      FULL      NONE
2:17     VR-Default  E      R      OFF 40000      FULL      NONE
2:18     VR-Default  E      NP     OFF 10000      FULL      NONE
2:19     VR-Default  E      NP     OFF 10000      FULL      NONE
=====

```

```

> indicates Port Display Name truncated past 8 characters
Link State: A-Active, R-Ready, NP-Port not present, L-Loopback
Port State: D-Disabled, E-Enabled
Media: !/$/%/* - use "show port info detail" for more information

```

```

# show port 2:1 info detail
Port:      2:1
Virtual-router: VR-Default
Type:      Q+SR4 ($ - Restricted Transceiver)
Random Early drop:      Unsupported
Admin state:      Enabled with 40G full-duplex
Link State:      Ready (local fault)
Link Ups:      1      Last: Wed Sep 05 04:38:19 2012
Link Downs:    1      Last: Wed Sep 05 20:35:04 2012

```

```

# show port 2:9 info detail
Port:      2:9

```

```

Virtual-router: VR-Default
Type:          Q+LR4 (% - Unrestricted Transceiver)
Random Early drop:      Unsupported
Admin state:    Enabled with 40G full-duplex
Link State:     Ready (local fault)
Link Ups:       1          Last: Wed Sep 05 04:38:19 2012
Link Downs:     1          Last: Wed Sep 05 20:35:04 2012
Virtual-router: VR-Default
Type:          SF+_LR (! - Unsupported Transceiver)
Random Early drop:      Unsupported
Admin state:    Enabled with 10G full-duplex
Link State:     Ready
Link Ups:       0          Last: --
Link Downs:     0          Last: --

```

The following example displays the current status of non-Extreme optical devices:

```

show port 3 transceiver information detail
Port : 3
Media Type          : Q+PSM4  (Licensed)
or
                               (Unlicensed)
or
                               (Restricted)
Vendor Name         : NOTEXT,INC
Part Number         : NOTEXT000EN-S001
Serial Number       : NE7M90299

Temp (Celsius)      : 0.00      Status          : Low Alarm
Low Warn Threshold  : 0.00      High Warn Threshold :
0.00
Low Alarm Threshold : 0.00      High Alarm Threshold : 0.00
Voltage AUX-1/Vcc (Volts) : 0.00      Status          : Low
Alarm
Low Warn Threshold  : 0.00      High Warn Threshold :
0.00
Low Alarm Threshold : 0.00      High Alarm Threshold :
0.00
Tx Power (dBm)      : -inf      Status          : Low
Alarm
Low Warn Threshold  : -inf      High Warn Threshold :
-inf
Low Alarm Threshold : -inf      High Alarm Threshold : -inf
Press <SPACE> to continue or <Q> to quit:

```



Note

In `show ports transceiver information` output, the Rx/Tx power values shown may be +/- 3dB from the actual value due to the limitation of SFP and the accuracy depends on SFP vendor. For accurate power measurement, it is recommended to use a power meter.

History

This command was first available in ExtremeXOS 12.4.

Support for the ExtremeSwitching switches was added in ExtremeXOS 12.5.

The tag value was added in ExtremeXOS 12.4.4.

Support for SFP and SFP+ optics was added in ExtremeXOS 12.5.3.

Show output was updated in 15.3.

Show output was updated to include the current status of non-Extreme optical devices in ExtremeXOS 15.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show ports transceiver information

```
show ports {port_list | tag tag} transceiver information
```

Description

Displays basic information about the optical transceiver.

Syntax Description

<i>port_list</i>	Specifies the port number(s).
<i>tag</i>	Specifies an 802.1Q or 802.1ad tag value.

Default

N/A.

Usage Guidelines

Digital Diagnostic Monitoring Interface (DDMI) provides critical system information about the installed optical modules. Use this command to monitor the condition of XFP, SFP, and SFP+ optical transceiver modules.

The tag value may be associated with either a VMAN or a VLAN.

If you try to execute this command on one of the ports in the port list that is non-compliant with DDMI, the following error message is displayed and the command does not go through:

```
Port 3:1 This command is not supported on this port. All ports and
transceiver of the ports requested in the command need to support DDMI.
```

If you try to execute this command on one of the ports in the port list on which the transceiver is non-compliant with DDMI, the following error message is displayed:

```
Port 3:1 This media/transceiver does not support enhanced digital
diagnostic monitoring interface (DDMI). All ports and transceiver of the
ports requested in the command need to support DDMI.
```

For more detailed information, use the `show ports transceiver information detail` command.

Example

The following display shows output for port 1:1-2:

```
# sh port 1:1-2 transceiver information
Port      Temp      TxPower  RxPower  TxBiasCurrent  Voltage-Aux1  Voltage-Aux2
(Celcius) (dBm)     (dBm)    (mA)       (Volts)        (Volts)
=====
1:1       30.60     -25.20   -18.70     0.40           5.09          5.07
1:2       30.60     -25.20   -18.70     0.40           5.09          N/A
=====
N/A indicates that the parameter is not applicable
to the optics connected to the port
```

The following display shows output for port 25:

```
# sh ports 25 transceiver information
Port      Temp      TxPower      RxPower
TxBiasCurrent  Voltage-Aux1  Voltage-Aux2
(mA)           (Celsius)    (dBm)        (dBm)
(Volts)                               (Volts)
=====
25            32.00        -3.35        -2.68
7.67          3.35        N/A
=====
```



Note

In `show ports transceiver information` output, the Rx/Tx power values shown may be +/- 3dB from the actual value due to the limitation of SFP and the accuracy depends on SFP vendor. For accurate power measurement, it is recommended to use a power meter.

History

This command was first available in ExtremeXOS 12.4.

Support for the ExtremeSwitching switches was added in ExtremeXOS 12.5.

The tag value was added in ExtremeXOS 12.4.4.

Support for SFP and SFP+ optics was added in ExtremeXOS 12.5.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show ports txerrors

```
show ports {port_list | stack-ports stacking-port-list} txerrors {no-
refresh | refresh}
```


Description

Displays real-time transmit error statistics. The switch automatically refreshes the output unless otherwise specified.

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
<i>stacking-port-list</i>	Specifies one or more stacking slot ports for display. Applies to SummitStack and ExtremeSwitching series switches only.
no-refresh	Specifies that auto-refresh is disabled. The output provides a real-time snapshot of the transmit errors at the time the command is issued. This setting is not saved.
refresh	Specifies a continuous refresh of output.

Default

The switch automatically refreshes the output.

Usage Guidelines

If you do not specify a port number or range of ports, error statistics are displayed for all ports.

If you do not specify the no-refresh parameter, the switch automatically refreshes the output (this is the default behavior).

If you specify the no-refresh parameter, the output provides a snapshot of the real-time transmit error statistics at the time you issue the command and displays the output in page-by-page mode (this was the default behavior in ExtremeXOS 11.2 and earlier). This setting is not saved; therefore, you must specify the no-refresh parameter each time you want a snapshot of the port transmit errors.

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

Collected Port Transmit Error Information

The switch collects the following port transmit error information:

- Port Number.
- Link State—The current state of the link. Options are:
 - Active (A)—The link is present at this port.
 - Ready (R)—The port is ready to accept a link.
 - Loopback (L)—The port is in Loopback mode.
- 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 Errors)—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 transmit frames that do not get completely transmitted because of buffer problems (FIFO underflow).
- Transmit Parity Frames (TX Parity)—The bit summation has a parity mismatch.

Port Monitoring Display Keys

For information about the available port monitoring display keys, see the [show ports statistics](#) command.

Example

The following command displays transmit error statistics for all ports on an ExtremeSwitching switch:

```
# show ports txerrors
```

The following is a truncated sample output from this command:

```

Port Tx Error Monitor
Port      Link      Tx      Tx      Tx      Tx      Tx      Tx
State     Coll     Late coll Deferred Errors  Lost      Parity
=====
1         R         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
9         R         0         0         0         0         0         0
10        R         0         0         0         0         0         0
11        R         0         0         0         0         0         0
12        R         0         0         0         0         0         0
13        R         0         0         0         0         0         0
14        R         0         0         0         0         0         0
15        R         0         0         0         0         0         0
16        R         0         0         0         0         0         0
17        R         0         0         0         0         0         0
=====
Link State: A-Active, R-Ready, NP-Port Not Present, L-Loopback
0->Clear Counters  U->page up  D->page down ESC->exitPort Tx Error

```

History

This command was first available in ExtremeXOS 10.1.

The link state information was updated to include NP-Port not present in ExtremeXOS 11.2.

Support for the auto-refresh functionality and the no-refresh parameter were added in ExtremeXOS 11.3. Auto-refresh continually updates the display. The no-refresh parameter takes a real-time snapshot of the display at the time you issue the command.

The **refresh** keyword was added in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show ports utilization

```
show ports {mgmt | port_list | tag tag | stack-ports stacking-port_list}  
          utilization {bandwidth | bytes | packet}
```

Description

Displays real-time port utilization information. The total utilization displays as real-time information, constantly refreshing, and the parameter displays show a snapshot of the activity on the port when you issue the command.

Syntax Description

mgmt	Specifies the management port.
<i>port_list</i>	Specifies one or more ports or slots and ports.
<i>tag</i>	Specifies an 802.1Q or 802.1ad tag value.
<i>stacking-port_list</i>	Specifies one or more stacking slots and ports.
bandwidth	Specifies port utilization as percentage of bandwidth.
bytes	Specifies port utilization in bytes per second.
packet	Specifies port utilization in packets per second.

Default

N/A.

Usage Guidelines

The software continuously monitors port utilization and calculates bandwidth as a function of each port's maximum link capacity.

The total utilization display presents real-time statistics. Use the *spacebar* to toggle the real-time displayed information for packets, bytes, and bandwidth in that order. When you use a parameter (packets, bytes, or bandwidth) with the command, the display for the specified type shows a snapshot per port when you issued the command. When the show ports utilization command is run with the bandwidth, bytes, or packets options, the command may need to be repeated a few times in order for the ExtremeXOS software to gather enough statistics to calculate appropriate values.

If you do not specify a port number, range of ports, or tag value, port utilization information is displayed for all ports.

The tag value may be associated with either a VMAN or a VLAN.

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 example shows sample output from the show ports utilization packets command:

```
Link Utilization Averages                               Mon Oct  6 22:38:25 2008
Port      Link      Rx          Peak Rx          Tx          Peak Tx
State     pkts/sec      pkts/sec      pkts/sec      pkts/sec
=====
1:1       A           47           191           0            0
1:2       A            0            0            0            0
2:1       R            0            0            0            0
2:2       R            0            0            0            0
3:1       R            0            0            0            0
3:2       R            0            0            0            0
4:1       R            0            0            0            0
4:2       R            0            0            0            0
5:1       R            0            0            0            0
5:2       R            0            0            0            0
6:1       R            0            0            0            0
6:2       R            0            0            0            0
7:1       R            0            0            0            0
7:2       R            0            0            0            0
=====
> indicates Port Display Name truncated past 8 characters
Link State: A-Active, R-Ready, NP-Port Not Present, L-Loopback
Spacebar->toggle screen U->page up  D->page down ESC->exit
```



Note

Use the *spacebar* to toggle this real-time display for all ports from packets to bytes to bandwidth, in that order.

The following example shows sample output from the show ports utilization bytes command:

```
Link Utilization Averages                               Mon Oct  6 22:39:22 2008
Port      Link      Rx          Peak Rx          Tx          Peak Tx
State     bytes/sec      bytes/sec      bytes/sec      bytes/sec
=====
1:1       A            0            0            0           63
1:2       A            0           63           63           63
2:1       R            0            0            0            0
2:2       R            0            0            0            0
3:1       R            0            0            0            0
3:2       R            0            0            0            0
4:1       R            0            0            0            0
4:2       R            0            0            0            0
```

```

5:1      R      0      0      0      0
5:2      R      0      0      0      0
6:1      R      0      0      0      0
6:2      R      0      0      0      0
7:1      R      0      0      0      0
7:2      R      0      0      0      0
=====
> indicates Port Display Name truncated past 8 characters
Link State: A-Active, R-Ready, NP-Port Not Present, L-Loopback
Spacebar->toggle screen U->page up D->page down ESC->exit

```

The following example shows sample output of the show ports utilization bandwidth command:

```

Link Utilization Averages
Port      Link      Link      Rx      Peak Rx      Mon Oct 6 22:39:46 2008
State     Speed    % bandwidth    % bandwidth    % bandwidth    Tx      Peak Tx
                                     % bandwidth    % bandwidth
=====
1:1      A      100      0.00      0.03      0.00      0.00
1:2      A      100      0.00      0.00      0.00      0.00
2:1      R      0       0.00      0.00      0.00      0.00
2:2      R      0       0.00      0.00      0.00      0.00
3:1      R      0       0.00      0.00      0.00      0.00
3:2      R      0       0.00      0.00      0.00      0.00
4:1      R      0       0.00      0.00      0.00      0.00
4:2      R      0       0.00      0.00      0.00      0.00
5:1      R      0       0.00      0.00      0.00      0.00
5:2      R      0       0.00      0.00      0.00      0.00
6:1      R      0       0.00      0.00      0.00      0.00
6:2      R      0       0.00      0.00      0.00      0.00
7:1      R      0       0.00      0.00      0.00      0.00
7:2      R      0       0.00      0.00      0.00      0.00
=====
> indicates Port Display Name truncated past 8 characters
Link State: A-Active, R-Ready, NP-Port Not Present, L-Loopback
Spacebar->toggle screen U->page up D->page down ESC->exit

```

History

This command was first available in ExtremeXOS 11.3.

The tag value was added in ExtremeXOS 12.4.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches. The stack-ports option is available only on SummitStack.

show ports vlan statistics

```
show ports {port_list} vlan statistics {no-refresh | refresh}
```

Description

Displays VLAN statistics at the port level.

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports. Can be one or more port numbers. May be in the form: 1, 2, 3-5, 1:*; 1:5, 1:6-1:8.
no-refresh	Specifies that there is no continuous refresh. The prompt comes back to the user after fetching statistics once.
refresh	Continuous refresh of output.

Default

N/A.

Usage Guidelines

This command is used in conjunction with the `configure ports [port_list|all] monitor vlanvlan_name {rx-only | tx-only}` command.

Example

The following command displays statistics for the ports 1-2 on node 5:

```
show ports 5:1-2 vlan stats
Displays the vlan statistics in a real time countinuous refresh mode or no-refresh mode.
# show ports 5:1-2 vlan stat
Port VLAN Statistics
Port      Vlan      Rx Frames      Rx Byte      Tx Frame      Tx Byte
Count      Count      Count      Count
=====
5:1      Default  318750522      20400046784      318750588      20400051672
5:2      Default  292811491      18739948736      292811975      18739980504
0->Clear Counters  U->page up  D->page down  ESC->exit
```

For ports that do not support transmit statistics, a '-' will be displayed. For ports that do not support transmit byte counters, a '-' will be displayed for that row and column. Similarly, configuration using rx-only or tx-only will result in the display of "-"s in the appropriate rows and columns.

History

This command was first available in ExtremeXOS 12.0.

The **refresh** keyword was added in ExtremeXOS 16.1.

Platform Availability

VLAN

show ports wred

```
show ports port_list wred ecn {no-refresh | refresh}
```

Description

Displays WRED and Explicit Congestion Notification (ECN) statistics for the specified ports or all ports.

Syntax Description

<i>port_list</i>	Specifies a list of slots and ports to display. Specify ports in the following formats: 3-5, 2:5, 2:6-2:8.
no-refresh	Specifies a static snapshot of data instead of the default dynamic display.
refresh	Continuous refresh of output.
ecn	Displays ECN statistics.

Default

N/A.

Usage Guidelines

If no port or port list is specified, this command displays the WRED statistics for all ports. If WRED is not configured on a port, the statistics for that port display as 0.

The drop counters in the display represent packets that were dropped based on the WRED congestion avoidance algorithm. The Green Pkt Drop column counts in-profile TCP and non-TCP packets that have been dropped. The Red Pkt Drop column counts out-of-profile TCP and non-TCP packets that have been dropped.



Note

The values in the Yellow Pkt Drop column are always 0 in this release because the yellow traffic color is not supported at this time.

Example

The following example displays the WRED statistics for port list 2:1-9:

```
# show ports 2:1-9 wred no-refresh
Port WRED Stats Monitor
=====
Port      Link      Green      Yellow      Red
State     Pkt Drop  Pkt Drop  Pkt Drop
=====
2:1       A         0          0          0
2:2       R         0          0          0
2:3       R         0          0          0
2:4       R         0          0          0
2:5       R         0          0          0
2:6       R         0          0          0
2:7       R         0          0          0
2:8       R         0          0          0
2:9       R         0          0          0
=====
Link State: A-Active, R-Ready, NP-Port Not Present L-Loopback
```

The following example shows ECN statistics for port 2:

```
# show ports 2 wred ecn no-refresh
Port WRED ECN Statistics Monitor
Port      Link      Packets
          State      Marked
=====
2         R         0
=====
Link State: A-Active, R-Ready, NP-Port Not Present L-Loopback
```

The following example shows ECN statistics for all ports:

```
# show ports wred ecn no-refresh
Port WRED ECN Statistics
Port      Link      Packets
          State      Marked
=====
1         A         0
2         R         0
3         A         0
4         A         0
5         R         0
6         R         0
7         R         0
8         R         0
9         R         0
10        R         0
11        R         0
12        R         0
13        R         0
14        R         0
15        R         0
16        R         0
17        R         0
18        R         0
19        R         0
20        R         0
21        R         0
22        R         0
23        R         0
24        R         0
25        R         0
26        R         0
27        R         0
28        R         0
29        R         0
30        R         0
31        R         0
32        R         0
33        NP         0
34        NP         0
=====
Link State: A-Active, R-Ready, NP-Port Not Present L-Loopback
```

History

This command was first available in ExtremeXOS 12.7.

The **refresh** keyword was added in ExtremeXOS 16.1.

The **ecn** keyword was added in ExtremeXOS 22.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show power

```
show power {ps_num} {detail}
```

Description

Displays the current status of the installed power supplies.

Syntax Description

<i>ps_num</i>	Specifies the slot number of the installed power supply.
detail	The detail option is reserved for future use.

Default

N/A.

Usage Guidelines

Use this command to view detailed information about the health of the power supplies.

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

The switch collects the following power supply information:

- State—Indicates the current state of the power supply. Options are:
 - Empty—There is no power supply installed.
 - Power Failed—The power supply has failed.
 - Powered Off—The power supply is off.
 - Powered On—The power supply is on and working normally.
- Disabled for net power gain—Indicates that the power supply is disabled in order to maximize the total available system power.
- Configured ON—Indicates that the user requested to enable a disabled power supply regardless of the affect on the total available system power.
- Configured ON when present—Indicates that the power supply slot is currently empty, but the user requested to enable the power supply regardless of the affect on the total available system power.
- Internal Power Supply (PowerSupply 1 information)—The ExtremeSwitching switches come with one power supply pre-installed at the factory. The ExtremeSwitching power supply is not user-replaceable; therefore, the part information display indicates internal power supply.

- External Power Supply (PowerSupply 2-4 information)—Displays information about the optional External Power System (EPS) that allows you to add a redundant power supply to the ExtremeSwitching series switches to protect against a power supply failure.

The ExtremeSwitching series switches include Power Usage, which is only an estimate for the input power consumed. SummitStack displays the supplies associated to each active node that is present in a slot. The supplies are represented with flags that describe whether the supply is providing power, has failed, or is providing no power, or if the supply has had its 48v power output automatically turned off because two or three external power supplies are available. For more information, see [show power \(Stack Nodes Only\)](#).

In ExtremeXOS 10.1 and earlier, use the `show power supplies {detail}` command to view detailed health information about the power supplies.

In ExtremeXOS 22.5 and later, the **detail** option shows external power supply fan airflow direction. This is only applicable for switches that support such power supplies:

- ExtremeSwitching 5320—all models
- ExtremeSwitching 5420—all models
- ExtremeSwitching 5520—all models
- ExtremeSwitching 5720—all models

Bridge Port Extenders (BPEs)

Power supply information for bridge port extenders (BPEs) does not appear in this command. Use the `show vpex bpe {slot slot_num} {environment}` command instead.

Example

The following command displays the status of the power supplies installed:

```
# show power
```

The following sample output assumes that you have not installed an EPS:

```
PowerSupply 1 information:
State:           Powered On
PartInfo:        Internal Power Supply
PowerSupply 2 information:
State:           Empty
```

The following sample output assumes that you have installed an EPS:

```
PowerSupply 1 information:
State:           Powered On
PartInfo:        Internal Power Supply
PowerSupply 2 information:
State:           Powered On
PartInfo:        External Power Supply
```

The following sample output assumes a connection to an external EPS-C with three EPS-600LS modules installed:

```
# show power
PowerSupply 1 information:
```

```

State:      Powered Off
PartInfo:   Internal Power Supply
Input:      0.00 V AC
PowerSupply 2 information:
State:      Powered On
PartInfo:   External Power Supply
Input:      0.00 V AC
PowerSupply 3 information:
State:      Powered On
PartInfo:   External Power Supply
Input:      0.00 V AC
PowerSupply 4 information:
State:      Powered On
PartInfo:   External Power Supply
Input:      0.00 V AC

```

For stacking systems, the power detail display is enhanced as follows:

```

# show power
PSU-1 or PSU-2 or
Internal External External External Power
Slots Type PSU PSU PSU PSU Usage
-----
Slot-1 P - - - 113.88 W
Slot-2
Slot-3
Slot-4
Slot-5
Slot-6
Slot-7
Slot-8
Flags : (P) Power available, (F) Failed or no power,
(O) 48V powered off when 2 or 3 external PSUs are powered on,
(-) Empty
System Power Usage : 120 W
* Slot-1 Stack.2 # show power detail
Slot-1 PowerSupply 1 information:
State : Powered On
PartInfo : PSSF751301A- 1022A-40459 800382-00-01
Power Usage : 113.88W
Output 1 : 18.63 V, 4.37 A
Output 2 : 3 V, 3 A
Slot-1 PowerSupply 2 information:
State : Empty
System Power Usage : 120 W

```

The following example shows power detail:

```

# show power detail

PowerSupply 1 information:
State : Powered On
PartInfo : Internal PSU-1 1402E-41506 800386-00-06
Input : 121.06 V AC
Output 1 : 12.01 V, 4.12 A (12V/300W Max)
Power Usage : 66.86 W
AirFlow Direction : Front To Back

```

History

This command was first available in an ExtremeXOS 10.1.

The syntax for this command was modified in ExtremeXOS 11.0 from `show powersupplies` to `show power {ps_num} {detail}`.

Power supply fan airflow direction information was added in ExtremeXOS 22.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show power (Stack Nodes Only)

```
show power
```

Description

Displays the number of power modules present and providing power in each slot.

Syntax Description

This command has no arguments or variables.

Default

Default value

Usage Guidelines

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches. However, it produces completely different output on a stack. The following table describes the flags that appear when this command is executed on an active node.

Table 43: Flag Descriptions for the show power Command

Power Supply	Flag	Meaning
Internal	F	Failed or no Power.
Internal	P	Power available.
Internal	O	48V powered off (48p Summits only).
External (non 48P)	-	Empty.
External (non 48P)	F	Failed or no power.
External (non 48P)	P	Power available.
External (non 48P)		Power supply can never occupy this position.
External (48P only)	-	Empty or no power to all PSUs present.
External (48P only)	F	Failed or no power (at least 1 PSU has power).
External (48P only)	P	Power available.

All ExtremeSwitching series switches accept an external power chassis that holds only one power supply. For ExtremeSwitching series switches, the External PSU columns are left blank.

For slots without active nodes, the slot number appears and the remainder of the row is blank.

Example

The following are sample displays for this command:

```
# show power
Internal  External  External  External
Slots    Type      PSU       PSU       PSU       PSU
-----
Slot-1 X          F         P
Slot-2 X          P         -
Slot-3 X          -         P
Slot-4 X          P         P
Slot-5 X          P         -
Slot-6 X          -         P
Slot-7 X          P         P         F         -
Slot-8 X          P         -
Flags : (P) Power available, (F) Failed or no power,
(O) 48V powered off when 2 or 3 external PSU are powered on,
(-) Empty
Slot-2 Stack.41 #
Slot-2 Stack250.4 # show power
Internal  External  External  External
Slots    Type      PSU       PSU       PSU       PSU
-----
Slot-1 X          F         P
Slot-2 X          P         -
Slot-3 X          P         -
Slot-4 X          P         P
Slot-5 X          P         -
Slot-6 X          O         P         P         F
Slot-7
Slot-8 X          P         -
Flags : (P) Power available, (F) Failed or no power,
(O) 48V powered off when 2 or 3 external PSUs are powered on,
(-) Empty
Slot-2 StackX #
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches. However, the output described in this section is available only on nodes configured for the SummitStack feature.

show private-vlan

```
show private-vlan
```

Description

Displays information about all the PVLANS on the switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

If the PVLAN is incomplete because it does not have a network or any subscriber VLAN configured, [INCOMPLETE] appears next to the PVLAN name.

Example

The following example output displays all the PVLANS on the switch:

```
# show private-vlan
-----
Name          VID  Protocol Addr      Flags          Proto  Ports  Virtual
Active router
/Total
-----
Engineering
Network VLAN:
-Engr1        10  -----
ANY          4 /5  VR-Default
Non-Isolated Subscriber VLAN:
-ni1          400 -----
ANY          1 /1  VR-Default
-ni2          401 -----
ANY          1 /1  VR-Default
Isolated Subscriber VLAN:
-il           500 -----
ANY          1 /1  VR-Default
Ops
Network VLAN:
-Ops          20  -----
ANY          2 /2  VR-Default
Non-Isolated Subscriber VLAN:
-OpsNi1       901 -----
ANY          1 /1  VR-Default
-OpsNi2       902 -----
ANY          1 /1  VR-Default
-OpsNi3       903 -----
ANY          1 /1  VR-Default
-OpsNi4       904 -----
ANY          1 /1  VR-Default
Isolated Subscriber VLAN:
-OpsI0        600 -----
ANY          1 /1  VR-Default
-OpsI1        601 -----
ANY          1 /1  VR-Default
-OpsI2        602 -----
ANY          1 /1  VR-Default
-OpsI3        603 -----
ANY          1 /1  VR-Default
-OpsI4        604 -----
ANY          1 /1  VR-Default
Sales [INCOMPLETE]
Network VLAN:
-NONE
Non-Isolated Subscriber VLAN:
-SalesNi1     701 -----
ANY          1 /1  VR-Default
-SalesNi2     702 -----
ANY          1 /1  VR-Default
Isolated Subscriber VLAN:
-SalesI0      800 -----
ANY          1 /1  VR-Default
-----
```

```

Flags : (C) EAPS Control vlan, (d) NetLogin Dynamically created VLAN,
(D) VLAN Admin Disabled, (E) ESRP Enabled, (f) IP Forwarding Enabled,
(i) ISIS Enabled, (I) IP Forwarding lpm-routing Enabled, (L) Loopback Enabled,
(l) MPLS Enabled, (m) IPmc Forwarding Enabled, (n) IP Multinetting Enabled,
(N) Network Login vlan, (o) OSPF Enabled, (p) PIM Enabled,
(P) EAPS protected vlan, (r) RIP Enabled,
(T) Member of STP Domain, (V) VPLS Enabled, (v) VRRP Enabled
Total number of PVLAN(s) : 3

```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms that support the Private VLAN feature. For features and the platforms that support them, see the [Switch Engine 32.3 Feature License Requirements](#) document.

show private-vlan name

```
show {private-vlan} name
```

Description

Displays information about the specified PVLAN.

Syntax Description

<i>name</i>	Specifies the name of the PVLAN to display.
-------------	---

Default

N/A.

Usage Guidelines

If the PVLAN is incomplete because it does not have a network or any subscriber VLAN configured, [INCOMPLETE] appears next to the PVLAN name.

Example

The following example output displays information for the companyx PVLAN:

```

# show private-vlan "Engineering"
-----
Name          VID  Protocol Addr      Flags                Proto  Ports  Virtual
Active router
/Total
-----
Engineering
Network VLAN:
-Engr1        10  -----            ANY      4 /5   VR-Default

```

```

Non-Isolated Subscriber VLAN:
-ni1      400 ----- ANY    1 /1  VR-Default
-ni2      401 ----- ANY    1 /1  VR-Default
Isolated Subscriber VLAN:
-il       500 ----- ANY    1 /1  VR-Default
-----
Flags : (C) EAPS Control vlan, (d) NetLogin Dynamically created VLAN,
(D) VLAN Admin Disabled, (E) ESRP Enabled, (f) IP Forwarding Enabled,
(i) ISIS Enabled, (I) IP Forwarding lpm-routing Enabled, (L) Loopback Enabled,
(l) MPLS Enabled, (m) IPmc Forwarding Enabled, (n) IP Multinetting Enabled,
(N) Network Login vlan, (o) OSPF Enabled, (p) PIM Enabled,
(P) EAPS protected vlan, (r) RIP Enabled,
(T) Member of STP Domain, (V) VPLS Enabled, (v) VRRP Enabled

```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on all platforms that support the Private VLAN feature. For features and the platforms that support them, see the [Switch Engine 32.3 Feature License Requirements](#) document.

show process

```
show process {name} {detail} {description} {slot slotid}
```

Description

Displays the status of the "Vital" processes.

Syntax Description

<i>name</i>	Specifies the name of the process.
detail	Specifies more detailed process information.
description	Describes the name of all of the processes or the specified process running on the switch.
<i>slotid</i>	Specifies the slot number of a node in a stack topology.

Default

N/A.

Usage Guidelines

The NOS process manager monitors all of the "Vital" processes. The process manager also ensures that only version-compatible processes are started.

Using this command without the optional keywords displays summary process information. When you specify the slot keyword, summary information is displayed for that particular slot only.

The `show process` and `show process slot slotid` commands display the following information in a tabular format:

- Process Name—The name of the process.
- Version—The version number of the process. Options are:
 - Version number—A series of numbers that identify the version number of the process. This is helpful to ensure that you have version-compatible processes and if you experience a problem.
 - Not Started—The process has not been started. This can be caused by not having the appropriate license or for not starting the process.
- Restart—The number of times the process has been restarted. This number increments by one each time a process stops and restarts.
- State—The current state of the process. Options are:
 - No License—The process requires a license level that you do not have. For example, you have not upgraded to that license, or the license is not available for your platform.
 - Ready—The process is running.
 - Stopped—The process has been stopped.
- Start Time—The current start time of the process. Options are:
 - Day/Month/Date/Time/Year—The date and time the process began. When a process terminates and restarts, the start time is also updated.
 - Not Started—The process has not been started. This can be caused by not having the appropriate license or for not starting the process.
- Group—The process control group to which the process belongs to. Options are:
 - Vital—The process belongs to “Vital” process control group.
 - Other—The process belongs to “Other” process control group.
 - Kernel—The process belongs to “root” process control group.

When you specify the **detail** keyword, more specific and detailed process information is displayed.

The `show process detail` and `show process slot slotid` commands display the following information in a multi-tabular format:

- Detailed process information
- Memory usage configurations
- Recovery policies
- Process statistics
- Resource usage

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

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

You may find it useful to capture the process information under normal operating conditions to establish a baseline. By having a baseline, if you experience a problem, you and your technical support representative can more easily identify the problem.

SummitStack Only

When you run the command with out any parameters:

- From the stack manager or backup node, the stack displays the status of the "Vital" processes running on the master node and the back-up node in the Active Topology.
- From a standby node, the stack displays the status of the "Vital" processes running on the standby node and the master node in the Active Topology.

Example

The following is sample output from an ExtremeSwitching switch:

# show process									
Process Name	Version	Restart	State	Start Time				Group	
aaa	3.0.0.4	0	Ready	Sat	Dec	11	22:42:28	2021	Vital
acl	3.0.0.2	0	Ready	Sat	Dec	11	22:42:28	2021	Vital
bfd	1.0.0.1	0	Ready	Sat	Dec	11	22:42:28	2021	Vital
bgp	4.0.0.2	0	Ready	Sat	Dec	11	22:42:28	2021	Vital
brm	1.0.0.0	0	Ready	Sat	Dec	11	22:42:29	2021	Vital
cfgmgr	3.0.0.21	0	Ready	Sat	Dec	11	22:42:28	2021	Vital
cli	3.0.0.22	0	Ready	Sat	Dec	11	22:42:28	2021	Vital
devmgr	3.0.0.2	0	Ready	Sat	Dec	11	22:42:28	2021	Vital
dirser	3.0.0.2	0	Ready	Sat	Dec	11	22:42:27	2021	Vital
dosprotect	3.0.0.1	0	Ready	Sat	Dec	11	22:42:28	2021	Vital
dotlag	1.0.0.1	0	Ready	Sat	Dec	11	22:42:29	2021	Vital
eaps	3.0.0.8	0	Ready	Sat	Dec	11	22:42:28	2021	Vital
edp	3.0.0.2	0	Ready	Sat	Dec	11	22:42:28	2021	Vital
elrp	3.0.0.1	0	Ready	Sat	Dec	11	22:42:28	2021	Vital
elsm	3.0.0.1	0	Ready	Sat	Dec	11	22:42:28	2021	Vital
ems	3.0.0.2	0	Ready	Sat	Dec	11	22:42:28	2021	Vital
epm	3.0.0.4	0	Ready	Sat	Dec	11	22:42:26	2021	Vital
erps	1.0.0.0	0	Ready	Sat	Dec	11	22:42:29	2021	Vital
esrp	3.0.0.4	0	Ready	Sat	Dec	11	22:42:28	2021	Vital
ethoam	1.0.0.1	0	Ready	Sat	Dec	11	22:42:28	2021	Vital
etmon	1.0.0.1	0	Ready	Sat	Dec	11	22:42:28	2021	Vital
exacl	3.0.0.2	0	Ready	Sat	Dec	11	22:42:27	2021	Kernel
exdhcpsnoop	1.0.0.1	0	Ready	Sat	Dec	11	22:42:27	2021	Kernel
exdos	3.0.0.2	0	Ready	Sat	Dec	11	22:42:26	2021	Kernel
exfib	1.0.0.2	0	Ready	Sat	Dec	11	22:42:27	2021	Kernel
exfipSnoop	1.0.0.0	0	Ready	Sat	Dec	11	22:42:27	2021	Kernel
exosmc	3.0.0.2	0	Ready	Sat	Dec	11	22:42:27	2021	Kernel
exosq	3.0.0.2	0	Ready	Sat	Dec	11	22:42:26	2021	Kernel
exsflow	1.0.0.2	0	Ready	Sat	Dec	11	22:42:27	2021	Kernel
exsnoop	3.0.0.2	0	Ready	Sat	Dec	11	22:42:27	2021	Kernel
exsshd	6.5.1.69	0	Ready	Sat	Dec	11	22:42:29	2021	Other
exvlan	3.0.0.2	0	Ready	Sat	Dec	11	22:42:27	2021	Kernel
fcoe	1.0.0.0	0	Ready	Sat	Dec	11	22:42:29	2021	Vital
fdb	7.1.0.0	0	Ready	Sat	Dec	11	22:42:28	2021	Vital
gptp	1.0.0.0	0	Ready	Sat	Dec	11	22:42:29	2021	Vital
hal	3.0.0.2	0	Ready	Sat	Dec	11	22:42:28	2021	Vital
hclag	1.0.0.0	0	Ready	Sat	Dec	11	22:42:29	2021	Vital
idMgr	1.0.1.1	0	Ready	Sat	Dec	11	22:42:29	2021	Vital
ipSecurity	1.0.0.1	0	Ready	Sat	Dec	11	22:42:29	2021	Vital
ipfix	3.0.0.2	0	Ready	Sat	Dec	11	22:42:28	2021	Vital
isis	1.0.0.2	0	Ready	Sat	Dec	11	22:42:29	2021	Vital
ismb	1.0.0.0	0	Ready	Sat	Dec	11	22:42:29	2021	Vital
lacp	3.0.0.1	0	Ready	Sat	Dec	11	22:42:28	2021	Vital
lldp	1.2.0.0	0	Ready	Sat	Dec	11	22:42:28	2021	Vital
mcmgr	4.0.0.2	0	Ready	Sat	Dec	11	22:42:28	2021	Vital
mpls	Not Started	0	No license	Not	Started				Vital

mrp	1.0.0.0	0	Ready	Sat Dec 11 22:42:29 2021	Vital
msdp	1.0.0.2	0	Ready	Sat Dec 11 22:42:28 2021	Vital
netLogin	2.1.0.1	0	Ready	Sat Dec 11 22:42:28 2021	Vital
netTools	3.0.0.2	0	Ready	Sat Dec 11 22:42:28 2021	Vital
nettx	3.0.0.2	0	Ready	Sat Dec 11 22:42:26 2021	Kernel
nodealias	1.0.0.1	0	Ready	Sat Dec 11 22:42:29 2021	Vital
nodealias_snoop	1.0.0.1	0	Ready	Sat Dec 11 22:42:27 2021	Kernel
nodemgr	3.0.0.2	0	Ready	Sat Dec 11 22:42:28 2021	Vital
ntp	4.2.6.3	0	Ready	Sat Dec 11 22:42:29 2021	Vital
ospf	3.0.0.3	0	Ready	Sat Dec 11 22:42:28 2021	Vital
ospfv3	3.0.0.2	0	Ready	Sat Dec 11 22:42:28 2021	Vital
otm	1.0.0.1	0	Ready	Sat Dec 11 22:42:29 2021	Vital
pim	3.0.0.2	0	Ready	Sat Dec 11 22:42:28 2021	Vital
polMgr	3.0.0.2	0	Ready	Sat Dec 11 22:42:28 2021	Vital
policy	1.0.0.1	0	Ready	Sat Dec 11 22:42:29 2021	Vital
pwmib	1.0.0.0	0	Ready	Sat Dec 11 22:42:28 2021	Vital
rip	3.0.0.2	0	Ready	Sat Dec 11 22:42:28 2021	Vital
ripng	3.0.0.1	0	Ready	Sat Dec 11 22:42:28 2021	Vital
rtmgr	4.0.0.2	0	Ready	Sat Dec 11 22:42:28 2021	Vital
snmpMaster	4.0.0.2	0	Ready	Sat Dec 11 22:42:28 2021	Vital
snmpSubagent	3.0.0.2	0	Ready	Sat Dec 11 22:42:28 2021	Vital
stp	3.0.4.4	0	Ready	Sat Dec 11 22:42:28 2021	Vital
techSupport	1.0.0.0	0	Ready	Sat Dec 11 22:42:28 2021	Vital
telnetd	3.0.0.2	0	Ready	Sat Dec 11 22:42:28 2021	Other
tftpd	3.0.0.2	0	Ready	Sat Dec 11 22:42:28 2021	Vital
thttpd	1.0.0.0	0	Ready	Sat Dec 11 22:42:29 2021	Other
twamp	1.0.0.1	0	Ready	Sat Dec 11 22:42:29 2021	Vital
upm	1.0.0.1	0	Ready	Sat Dec 11 22:42:29 2021	Vital
vlan	4.1.0.3	0	Ready	Sat Dec 11 22:42:28 2021	Vital
vmt	1.0.1.1	0	Ready	Sat Dec 11 22:42:29 2021	Vital
vrrp	3.0.0.5	0	Ready	Sat Dec 11 22:42:28 2021	Vital
vsm	1.0.0.2	0	Ready	Sat Dec 11 22:42:29 2021	Vital
xmlc	1.0.1.0	0	Ready	Sat Dec 11 22:42:29 2021	Vital
xmld	1.0.0.0	0	Ready	Sat Dec 11 22:42:28 2021	Vital

The following example describes the name of all of the processes running on the switch:

```
show process description
```

The following is sample output from this command:

```
# show process description
Process Name      Description
-----
aaa               Authentication, Authorization, and Accounting Server
acl               Access Control List Manager
bfd               IETF Bidirectional Forwarding Detection
bgp               Border Gateway Protocol
brm               Bandwidth Resource Manager
cfgmgr            Configuration Manager
cli               Cli Manager
devmgr            Device Manager
dirser            Directory Services
dosprotect        protects against Denial of Service attacks
dotlag            IEEE 802.1ag; Connectivity Fault Management
eaps              Ethernet Automatic Protection Switching
edp               Extreme Discovery Protocol
elrp              Extreme Loop Recovery Protocol
elsm              Extreme Link State Monitor
ems               Event Management System server application
epm               Extreme Process Manager
esrp              Extreme Standby Routing Protocol
ethoam            Ethernet OAM
etmon             Traffic monitoring and sampling utility
```

exacl	Access Control List Module
exdhcpsnoop	DHCP snooping module
exdos	Detection of potential Denial of Service attacks module
exfib	Routing interface to manage missing routes in ASIC
exosipv6	IPv6 Custom Interface Module
exosmc	Multicast Forwarding Module
exosnvram	Interface to non-volatile RAM
exosq	EXOS Queue Module
exsflow	Sflow interface to gather sflow samples
exsnnoop	IGMP/MLD Snooping Module
exvlan	Layer 2 configuration module
fdb	Forwarding Data Base Manager
hal	Hardware Abstraction Layer
hclag	Health Check LAG
idMgr	Identity Manager
ipSecurity	IP Security
ipfix	IPFIX Traffic monitoring utility
isis	Intermediate System to Intermediate System Route Protocol
lacp	Link Aggregation Control Protocol
lldp	802.1AB; Station and Media Access Control Connectivity Discover
mcmgr	Multicast Cache Manager
mpls	Multi-Protocol Label Switching
msdp	Multicast Source Discovery Protocol
msgsrv	Message Server
netLogin	Network Login includes MAC, Web-Based and 802.1X authentication
netTools	Network Tools set includes ping/tracert/bootprelay/dhcp/dns/sn
nettx	Layer 2 forwarding engine module
nodemgr	Fault Tolerance Manager
ospf	Open Shortest Path First Routing Protocol
ospfv3	Open Shortest Path First Routing Protocol for IPv6
pim	Protocol Independent Multicast
poe	Power Over Ethernet Manager
polMgr	Policy Manager
rip	Routing Information Protocol
ripng	Routing Information Protocol for IPv6
rtmgr	Route Table Manager
snmpMaster	Simple Network Management Protocol - Master agent
snmpSubagent	Simple Network Management Protocol - Subagent
stp	Spanning Tree Protocol
synce	Synchronous Ethernet
telnetd	Telnet server
tftpd	Tftp server
thttpd	Web Server
upm	Universal Port Manager
vlan	VLAN Manager - L2 Switching application
vmt	Virtual Machine Tracking
vrrp	Virtual Router Redundancy Protocol (RFC 3768)
vsm	Virtual Switch Manager
xmlc	XML Client Manager
xmld	XML server

The following example shows the truncated output for the command on a stack:

# show process									
Card	Process Name	Version	Restart	State	Start Time				Group

Slot-1	aaa	3.0.0.4	0	Ready	Mon Feb 6	16:01:24	2017	Vital	
Slot-1	acl	3.0.0.2	0	Ready	Mon Feb 6	16:01:24	2017	Vital	
Slot-1	bfd	1.0.0.1	0	Ready	Mon Feb 6	16:01:24	2017	Vital	
Slot-1	bgp	4.0.0.2	0	Ready	Mon Feb 6	16:01:24	2017	Vital	
Slot-1	brm	1.0.0.0	0	Ready	Mon Feb 6	16:01:24	2017	Vital	
Slot-1	cfgmgr	3.0.0.21	0	Ready	Mon Feb 6	16:01:24	2017	Vital	
Slot-1	cli	3.0.0.22	0	Ready	Mon Feb 6	16:01:24	2017	Vital	
Slot-1	devmgr	3.0.0.2	0	Ready	Mon Feb 6	16:01:24	2017	Vital	

```

Slot-1 dirser      3.0.0.2    0    Ready    Mon Feb 6 16:01:23 2017 Vital
Slot-1 dosprotect  3.0.0.1    0    Ready    Mon Feb 6 16:01:24 2017 Vital
Slot-1 dotlag      1.0.0.1    0    Ready    Mon Feb 6 16:01:24 2017 Vital
Slot-1 eaps        3.0.0.8    0    Ready    Mon Feb 6 16:01:24 2017 Vital
Slot-1 edp         3.0.0.2    0    Ready    Mon Feb 6 16:01:24 2017 Vital
Slot-1 elrp        3.0.0.1    0    Ready    Mon Feb 6 16:01:24 2017 Vital
Slot-1 elsm        3.0.0.1    0    Ready    Mon Feb 6 16:01:24 2017 Vital
Slot-1 ems         3.0.0.2    0    Ready    Mon Feb 6 16:01:24 2017 Vital
Slot-1 epm         3.0.0.4    0    Ready    Mon Feb 6 16:01:21 2017 Vital
Slot-1 erps        1.0.0.0    0    Ready    Mon Feb 6 16:01:24 2017 Vital
Slot-1 esrp        3.0.0.4    0    Ready    Mon Feb 6 16:01:24 2017 Vital
Slot-1 ethoam      1.0.0.1    0    Ready    Mon Feb 6 16:01:24 2017 Vital
Slot-1 etmon       1.0.0.1    0    Ready    Mon Feb 6 16:01:24 2017 Vital
.....
.....
Slot-2 aaa         3.0.0.4    0    Ready    Mon Feb 6 16:01:24 2017 Vital
Slot-2 acl         3.0.0.2    0    Ready    Mon Feb 6 16:01:24 2017 Vital
Slot-2 bfd         1.0.0.1    0    Ready    Mon Feb 6 16:01:24 2017 Vital
Slot-2 bgp         4.0.0.2    0    Ready    Mon Feb 6 16:01:24 2017 Vital
Slot-2 brm         1.0.0.0    0    Ready    Mon Feb 6 16:01:24 2017 Vital
Slot-2 cfgmgr      3.0.0.21   0    Ready    Mon Feb 6 16:01:24 2017 Vital
Slot-2 cli         3.0.0.22   0    Ready    Mon Feb 6 16:01:24 2017 Vital
Slot-2 devmgr      3.0.0.2    0    Ready    Mon Feb 6 16:01:24 2017 Vital
Slot-2 dirser      3.0.0.2    0    Ready    Mon Feb 6 16:01:22 2017 Vital
Slot-2 dosprotect  3.0.0.1    0    Ready    Mon Feb 6 16:01:24 2017 Vital
Slot-2 dotlag      1.0.0.1    0    Ready    Mon Feb 6 16:01:24 2017 Vital
Slot-2 eaps        3.0.0.8    0    Ready    Mon Feb 6 16:01:24 2017 Vital
Slot-2 edp         3.0.0.2    0    Ready    Mon Feb 6 16:01:24 2017 Vital
Slot-2 elrp        3.0.0.1    0    Ready    Mon Feb 6 16:01:24 2017 Vital
Slot-2 elsm        3.0.0.1    0    Ready    Mon Feb 6 16:01:24 2017 Vital
Slot-2 ems         3.0.0.2    0    Ready    Mon Feb 6 16:01:24 2017 Vital
Slot-2 epm         3.0.0.4    0    Ready    Mon Feb 6 16:01:21 2017 Vital
Slot-2 erps        1.0.0.0    0    Ready    Mon Feb 6 16:01:24 2017 Vital
Slot-2 esrp        3.0.0.4    0    Ready    Mon Feb 6 16:01:24 2017 Vital
Slot-2 ethoam      1.0.0.1    0    Ready    Mon Feb 6 16:01:24 2017 Vital
Slot-2 etmon       1.0.0.1    0    Ready    Mon Feb 6 16:01:24 2017 Vital
.....
.....

```

History

This command was first available in an ExtremeXOS 10.1.

The **description** keyword was added in ExtremeXOS 11.2.

Added in ExtremeXOS 15.7, the Version field will be overloaded to contain “User” if the process is a user created process.

Process group information was added in ExtremeXOS 22.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show process group

```
show process group {[vital | other]}
```

Description

This command shows the configured settings and statistics for the process groups.

Syntax Description

vital	Shows configured settings and statistics for the "Vital" process group.
other	Shows configured settings and statistics for the "Other" process group.

Default

If you do not select a process group ("Vital" or "Other"), information for both process groups appears.

Example

The following example shows information about both the "Vital" and "Other" process groups:

```
# show process group
  Number of processes          :   72
CPU
  Limit                        :   90 %
  Current utilization          :    1.2 %
  Maximum utilization          :   29.1 %
Memory
  Upper Limit                  :   95 %
  Current utilization          :    3.9 %
  Maximum utilization          :    3.9 %
  Number of processes killed by EPM :    0
Other:
  Number of processes          :    1
CPU
  Limit                        :   10 %
  Current utilization          :    0.0 %
  Maximum utilization          :    0.0 %
Memory
  Upper Limit                  :    5 %
  Current utilization          :    0.0 %
  Maximum utilization          :    0.0 %
  Number of processes killed by EPM :    0
```

History

This command was first available in ExtremeXOS 22.2.

The **exos** option was removed in ExtremeXOS 31.5.

The **vital** option was first available in ExtremeXOS 31.5

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show protocol

```
show protocol {filter} {filter_name} {detail}
```

Description

Displays protocol filter definitions and the complete protocol configuration.

Syntax Description

filter	Displays a protocol filter.
<i>name</i>	Displays a protocol filter name.
detail	Displays protocol information in detail.

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 command's output:

```
# show protocol
Protocol Filter Name   Protocol   Id       Destination   Field   Field   Field
                        Type       Value    Address      Offset  Value   Mask
-----
IP                     etype      0x0800
                       etype      0x0806
ANY                    ANY         0xffff
ipx                    etype      0x8137
IPv6                   etype      0x86dd
lacp                   etype      0x8809    01:80:C2:00:00:02>  14      01      FF
mpls                   etype      0x8847
appletalk              snap       0x809b
                       snap       0x80f3

> indicates that the value was truncated to the column size in the output.
Use the "detail" option to see the complete value.
```

The following example displays the `show protocol detail` command:

```
show protocol detail
Protocol Filter Name   : appletalk
  Protocol Id Type     : snap
  Protocol Id Value    : 0x809b
  Destination Address:
  Field Offset        :
  Field Value         :
  Field Mask          :
```

```

Protocol Id Type   : snap
Protocol Id Value  : 0x80f3
Destination Address:
Field Offset      :
Field Value       :
Field Mask        :

Protocol Filter Name : lacp
Protocol Id Type    : etype
Protocol Id Value   : 0x8809
Destination Address: 01:80:C2:00:00:02
Field Offset       : 14
Field Value        : 01
Field Mask         : FF

# show protocol filter "lacp" detail
Protocol Filter Name : lacp
Protocol Id Type     : etype
Protocol Id Value    : 0x8809
Destination Address: 01:80:C2:00:00:02
Field Offset        : 14
Field Value         : 01
Field Mask          : FF

```

The following example displays the `show protocol filter` command:

```

# show protocol filter
Protocol Name          Protocol Id  Destination  Field  Field  Field  Tagged
                        Type    Value      Address    Offset Value  Mask
-----
decnet                 etype  0x6003
                        etype  0x6004
myPvst                 snap   0x010b  01:00:0c:cc:cc:cd
netbios                llc    0xf0f0
                        llc    0xf0f1

```

History

This command was first available in ExtremeXOS 10.1.

The **filter** and **detail** keywords were added in ExtremeXOS 15.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show qosprofile

```

show qosprofile {ingress | egress} [ports [port_list | all | port_group]
                | NULL]

```

Description

Displays QoS information on the switch.

Syntax Description

ports	Specifies to show ports.
<i>port_list</i>	Specifies a list of ports.
<i>port_group</i>	Specifies the port group name.
all	Specifies all ports.
NULL	NULL.

Default

Displays egress QoS information for all ports.

Usage Guidelines

The displayed QoS profile information differs depending on the platform you are running on. The following section shows examples for different platforms.

Example

The display varies depending on your platform.

The following shows the information that appears when you omit the optional port parameter:

```
# show qosprofile
QP1   Weight = 1      Max Buffer Percent = 100
QP2   Weight = 1      Max Buffer Percent = 100
QP8   Weight = 1      Max Buffer Percent = 100
```

The following example shows how the display appears when the switch is configured for weighted-round-robin mode and some QoS profiles are configured for strict priority mode:

```
# show qosprofile
QP1   Weight = 1      Max Buffer Percent = 100
QP2   Weight = 1      Max Buffer Percent = 100
QP3   Weight = 1      Max Buffer Percent = 100
QP5   Strict-Priority Max Buffer Percent = 100
QP8   Strict-Priority Max Buffer Percent = 100
```

When you add the optional port parameter, the switch displays the following sample output:

```
Switch:6 # show qosprofile ports 1:1-2
Port: 1:1
QP1   MinBw =    20% MaxBw =    50% MaxBuf =   100%
QP8   MinBw =     0% MaxBw =   100% MaxBuf =  1000%
Port: 1:2
QP1   MinBw =     0% MaxBw =   100% MaxBuf =   100%
QP8   MinBw =     0% MaxBw =   100% MaxBuf =   100%
```



Note

This last sample output is not available on the XGS2 ports.

When you add the optional port group, the switch displays the following sample output:

```
# show qosprofile ports qosGroupA

Group: qosGroupA
Ports: 1-3
QP1  MinBw =    0% MaxBw      =   100% MaxBuf =   100% Weight =    1
QP3  MinBw =    0% MaxBw      =    34% MaxBuf =   100% Weight =    1
QP4  MinBw =    0% MaxBw      =    10% MaxBuf =   100% Weight =    1
QP5  MinBw =    0% MaxBw      =     1% MaxBuf =   100% Weight =    1
QP8  MinBw =    0% MaxBw      =   100% MaxBuf =   100% Weight =    1
```

History

This command was first available in ExtremeXOS 10.1.

The ingress information was added in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show qosscheduler

```
show qosscheduler {ports [ port_list | port_group | all ]}
```

Description

This command displays the scheduling algorithm that the switch uses to service [QoS](#) profiles for port lists or groups.

Syntax Description

qosscheduler	Get the QOS scheduling algorithm.
ports	Ports to display.
<i>port_list</i>	Port list.
<i>port_group</i>	Port group name.
all	All.

Default

N/A.

Usage Guidelines

Use this command to display the scheduling algorithm the switch uses to service QoS profiles for port lists or groups.

Example

```
# show qoscheduler ports "qosGroupA"

Group Name          : qosGroupA
Ports               : 1-3
Configured Scheduler : weighted-deficit-round-robin

# show qoscheduler ports 1-4,7
Port   Configured Scheduler
-----
1      weighted-round-robin
2      weighted-deficit-round-robin
3      strict-priority
4      strict-priority
7      weighted-round-robin
```

History

This command was first available in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show radius

```
show radius {mgmt-access | netlogin} {primary | secondary | index}
```

Description

Displays the current *RADIUS* client configuration and statistics.

Syntax Description

mgmt-access	Specifies configuration and statistics for the switch management RADIUS authentication server.
netlogin	Specifies configuration and statistics for the network login RADIUS authentication server.
primary	Primary server.
secondary	Secondary server.
<i>index</i>	RADIUS server index.

Default

N/A.

Usage Guidelines

If you do not specify a keyword, configuration details related to both management and network login are displayed. The output from this command displays the status of RADIUS and RADIUS accounting (enabled or disabled) and the primary and secondary servers for RADIUS and RADIUS accounting.

Use the **mgmt-access** keyword to display only RADIUS configuration details related to management access.

Use the **netlogin** keyword to only RADIUS configuration details related to network login.

Example

The following sample output displays the current RADIUS client configuration and statistics for both management and network login:

```
# show radius
RADIUS Default State:    enabled
RADIUS Default Timeout: 3 seconds
RADIUS Algorithm: standard
RADIUS Retries: 3
RADIUS dynamic-authorization: enabled
RADIUS TLS TCP Timeout: default
RADIUS TLS OCSP: on
Switch Management RADIUS: disabled
Switch Management RADIUS server connect time out: 3 seconds *
Switch Management RADIUS Accounting: disabled *
Switch Management RADIUS Accounting server connect time out: 3 seconds
Netlogin RADIUS: enabled
Netlogin RADIUS server connect time out: 3 seconds *
Netlogin RADIUS Accounting: disabled *
Netlogin RADIUS Accounting server connect time out: 3 seconds

RADIUS server          : 1 Status is Active
  host name            :
  IP address           : 10.51.1.150
  Server IP Port       : 1812
  Protocol              : UDP
  Client address       : 10.68.5.4 (VR-Mgmt)
  Retries               : 3 *
  Timeout              : 3 *
  Realm                : Netlogin
  shared secret        : #$Qzamzk10wQoU9jmTyFiEH1qT+Hp6+g==
Access Requests       : 2          Access Accepts       : 2
Access Rejects        : 0          Access Challenges  : 0
Access Retransmits    : 0          Client timeouts   : 0
Bad authenticators    : 0          Unknown types     : 0
Round Trip Time       : 221

RADIUS DynAuth server: 1 Status is Active
  host name            :
  IP address           : 10.51.1.150
  Server IP Port       : 3799
  Protocol              : UDP
  Client address       : 10.68.5.4 (VR-Mgmt)
  shared secret        : #$Mpu4FUpTNAtZ2cQAM/xAQI92SpD9vw==
  NAS-IP               : Require
CoA Requests          : 11          Disc Requests      : 5
CoA Accepts           : 2          Disc Accepts       : 3
CoA Rejects           : 7          Disc Rejects       : 1
CoA Dup Requests      : 0          Disc Dup Requests  : 0
```

```
CoA Bad Auths      : 0          Disc Bad Auths      : 0
CoA Packet Drops   : 2          Disc Packet Drops   : 1
```

The following command displays RADIUS TLS OSCP attributes (lines 9-12):

```
# show radius
RADIUS Default State:  enabled
RADIUS Default Timeout: 3 seconds
RADIUS Algorithm: standard
RADIUS Retries: 3
RADIUS dynamic-authorization: disabled
RADIUS TLS TCP Timeout: default
RADIUS TLS OSCP: on
RADIUS TLS OSCP Attributes:
  Nonce          : on
  Signer ocsf-nocheck : on
  Override Server URL : http://radiusocsp:2021
Switch Management RADIUS: disabled
Switch Management RADIUS server connect time out: 3 seconds *
Switch Management RADIUS Accounting: disabled *
Switch Management RADIUS Accounting server connect time out: 3 seconds
Netlogin RADIUS: enabled
Netlogin RADIUS server connect time out: 3 seconds *
Netlogin RADIUS Accounting: disabled *
Netlogin RADIUS Accounting server connect time out: 3 seconds

Primary Netlogin RADIUS server: Status is Active
  host name      :
  IP address     : 10.127.6.195
  Server IP Port : 1812
  Protocol       : UDP
  Client address : 10.127.6.85 (VR-Mgmt)
  Retries        : 3 *
  Timeout        : 3 *
  shared secret  : #$1HkCDc0zAm64sGwES6xVTN91clZEXQ==
Access Requests : 655          Access Accepts      : 655
Access Rejects  : 0            Access Challenges : 0
Access Retransmits: 0          Client timeouts    : 0
Bad authenticators: 0          Unknown types      : 0
Round Trip Time : 0

Legend: An asterisk (*) indicates a global value is in use.
```

History

This command was first available in ExtremeXOS 10.1.

The **mgmt-access** and **netlogin** keywords were added in ExtremeXOS 11.2.

The **primary** and **secondary** keywords, and *index* variable were added in ExtremeXOS 16.1.

This command was updated to show dynamic authorization status in ExtremeXOS 22.1.

This command was updated to show counters for dynamic authorization in ExtremeXOS 31.4.

RADIUS TLS OSCP attributes were added in ExtremeXOS 32.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show radius-accounting

```
show radius-accounting {mgmt-access | netlogin} {primary | secondary |  
                        index}
```

Description

Displays the current *RADIUS* accounting client configuration and statistics.

Syntax Description

mgmt-access	Specifies configuration and statistics for the switch management RADIUS accounting server.
netlogin	Specifies configuration and statistics for the network login RADIUS accounting server.
primary	Primary RADIUS accounting server.
secondary	Secondary RADIUS accounting server.
<i>index</i>	RADIUS accounting server index.

Default

N/A.

Usage Guidelines

If you do not specify a keyword, configuration details related to both management and network login are displayed. The output from this command displays information about the status and configuration of RADIUS accounting.

Use the **mgmt-access** keyword to display only RADIUS accounting configuration details related to management access.

Use the **netlogin** keyword to display only RADIUS accounting configuration details related to network login.

Example

The following sample output displays RADIUS accounting client configuration and statistics for both management and network login:

```
#show radius-accounting  
Radius Algorithm: standard  
Radius Retries: 3  
Switch Management Radius: enabled  
Switch Management Radius server connect time out: 3 seconds *  
Switch Management Radius Accounting: disabled *
```

```

Switch Management Radius Accounting server connect time out: 3 seconds *
Netlogin Radius: enabled
Netlogin Radius server connect time out: 3 seconds *
Netlogin Radius Accounting: disabled *
Netlogin Radius Accounting server connect time out: 3 seconds *

Primary Switch Management Radius server: Status: Active
    host name      :
    IP address     : 10.51.1.150
    Server IP Port : 1812
    Client address : 10.50.130.152 VR-Mgmt)
    Retries        : Not Set (Using Default Retry value)
    Timeout        : Not Set (Using Default Timeout value)
    shared secret  : @~|ioBnao`ox#s
Access Requests   : 0             Access Accepts      : 0
Access Rejects    : 0             Access Challenges : 0
Access Retransmits: 0             Client timeouts   : 0
Bad authenticators: 0             Unknown types      : 0
Round Trip Time   : 0

Radius server     : 1001 Status: Active
    host name      :
    IP address     : 10.51.1.150
    Server IP Port : 1812
    Client address : 10.50.130.152 VR-Mgmt)
    Retries        : Not Set (Using Default Retry value)
    Timeout        : Not Set (Using Default Timeout value)
    Realm          : Any
    shared secret  : @~|ioBnao`ox#s
Access Requests   : 11            Access Accepts      : 5
Access Rejects    : 1             Access Challenges : 0
Access Retransmits: 4             Client timeouts   : 5
Bad authenticators: 0             Unknown types      : 0
Round Trip Time   : 0

```

History

This command was first available in ExtremeXOS 10.1.

The `mgmt-access` and `netlogin` keywords were added in ExtremeXOS 11.2.

The **primary** and **secondary** keywords, and `index` variable were added in ExtremeXOS 16.1

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show radius dynamic-authorization

```
show radius dynamic-authorization index
```

Description

Displays the current *RADIUS* client dynamic-authorization status.

Syntax Description

<i>index</i>	RADIUS server index.
--------------	----------------------

Default

N/A.

Example

The following sample output displays the current RADIUS dynamic-authorization status:

```
# show radius dynamic-authorization
RADIUS dynamic-authorization: disabled
```

Example

The following sample output displays the NAS-IP Address requirement set to require (default).

```
# show radius dynamic-authorization
RADIUS dynamic-authorization: enabled

RADIUS DynAuth server: 1 Status is Active
  host name      :
  IP address     : 10.51.1.181
  Server IP Port : 3799
  Protocol       : UDP
  Client address : 10.50.97.80 (VR-Mgmt)
  shared secret  : #5LG11P/fwCjVyH1L1EqSeBrbGr/EXKQ==
  NAS-IP        : Require
```

History

This command was first available in ExtremeXOS 22.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show rip

```
show rip
```

Description

Displays *RIP* specific configuration.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays RIP specific configuration:

```
# show rip
```

The following is sample output from this command:

```
# show rip
RIP Routing      : Disabled      Operational status: Down
Split Horizon    : Enabled       Poison Reverse     : Enabled
Triggered Updates: Enabled       Aggregation        : Disabled
Update Interval  : 30             Route Timeout      : 180
Garbage Timeout  : 120            Router Alert       : Disabled
Originate Default: Disabled
Sys Import-Policy: None
Redistribute:
Protocol  Status   Cost Tag Policy
-----
Direct    Disabled 0    0    none
Static    Disabled 0    0    none
OSPFIntra Disabled 0    0    none
OSPFInter Disabled 0    0    none
OSPFExt1  Disabled 0    0    none
OSPFExt2  Disabled 0    0    none
E-BGP     Disabled 0    0    none
I-BGP     Disabled 0    0    none
ISISL1    Disabled 0    0    none
ISISL2    Disabled 0    0    none
ISISL1Ext Disabled 0    0    none
ISISL2Ext Disabled 0    0    none
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms with a Base license or higher as described in the [***Switch Engine 32.3 Feature License Requirements***](#) document.

show rip interface vlan

```
show rip interface vlan vlan_name
```

Description

Displays *RIP* specific statistics and configuration for a *VLAN* in detail.

Syntax Description

<i>vlan_name</i>	Specifies a VLAN name.
------------------	------------------------

Default

N/A.

Usage Guidelines

None.

Example

The following command displays RIP specific statistics for the VLAN accounting:

```
# show rip interface vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms with a Base license or higher as described in the *Switch Engine 32.3 Feature License Requirements* document.

show rip interface

show rip interface {**detail**}

Description

Displays *RIP*-specific configuration and statistics for all VLANs.

Syntax Description

detail	Specifies detailed display.
---------------	-----------------------------

Default

Show summary output for all interfaces.

Usage Guidelines

Summary includes the following information per interface:

- VLAN name.
- IP address and mask.
- interface status.
- packets transmitted.
- packets received.
- number of triggered updates.
- cost.

Detail includes the following per interface:

- VLAN name.
- IP address and mask.
- tx mode.
- rx mode.
- cost.
- peer information (for each peer).
 - age.
 - version.
 - received packets.
 - received updates.
 - received bad packets.
 - received bad routes.
- in policy.
- out policy.
- trusted gateway policy.
- packets transmitted.
- sent triggered updates.
- packets received.
- bad packets received.
- bad routes received.

Example

The following command displays the RIP configuration for all VLANs:

```
# show rip interface
```

The following is sample output from this command:

```
# show rip interface
VLAN      IP Address      Flags  Sent    Rcvd    Triggered Cost
Packets  Packets  Updates
Flags: (f) Interface Forwarding Enabled, (i) Interface RIP Enabled
(n) Multinetted VLAN, (r) Router RIP Enabled
```

The following command displays RIP-specific statistics for all VLANs:

```
# show rip interface detail
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms with a Base license or higher as described in the [*Switch Engine 32.3 Feature License Requirements*](#) document.

show rip memory

```
show rip memory {detail | memoryType}
```

Description

Displays *RIP* specific memory usage.

Syntax Description

detail	Displays detail information.
<i>memoryType</i>	Specifies the memory type usage to display.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays RIP specific memory for TEST:

```
# show rip memory test
```

The following is sample output from this command:

```
# show rip memory test
RIP Memory Information
-----
Bytes Allocated: 0          AllocFailed: 0
Current Memory Utilization Level: GREEN
Memory Utilization Statistics
-----
Memory Statistics for TEST
-----
```

Size	64	80	96	256	384	512	768	1024	2048	40
96	18432									

--										
Allocated		0	0	0	0	0	0	0	0	0
0	0									
AllocatedPeak		0	0	0	0	0	0	0	0	0
0	0									
AllocSuccess		0	0	0	0	0	0	0	0	0
0	0									
FreeSuccess		0	0	0	0	0	0	0	0	0
0	0									
AllocFail		0	0	0	0	0	0	0	0	0
0	0									
FreeFail		0	0	0	0	0	0	0	0	0
0	0									

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms with a Base license or higher as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

show rip routes

```
show rip routes {detail} {network ripNetworkPrefix}
```

Description

Displays routes advertised by [RIP](#).

Syntax Description

detail	Displays all available information from the RIP routing table.
<i>ripNetworkPrefix</i>	Specifies the route prefix for the routes to show.

Default

N/A.

Usage Guidelines

The routes displayed include all routes advertised by RIP, including routes exported from the system routing table and originated by other protocols, for example [BGP](#).

Example

The following command displays a summary of RIP specific routes for the networks 10.0.0.0/8:

```
# show rip routes network 10.0.0.0/8
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms with a Base license or higher as described in the [***Switch Engine 32.3 Feature License Requirements***](#) document.

show ripng

```
show ripng
```

Description

Displays *RIPng* global configuration and runtime information.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays RIPng global configuration and runtime information:

```
show ripng
```

The following is sample output from this command:

```
RIPng Routing      : Disabled
Aggregation       : Disabled
Update Interval   : 30           Route Timeout      : 180
Garbage Timeout   : 120
Originate Default : Disabled
Sys Import-Policy : None
Redistribute:
Protocol          Status    Cost Tag    Policy
```

Direct	Disabled	0	0	none	
Static	Disabled	0	0	none	
Ospf3-intra	Disabled	0	0	none	
Ospf3-inter	Disabled	0	0	none	
Ospf3-extern1	Disabled	0	0	none	
Ospf3-extern2	Disabled	0	0	none	
IsisL1	Disabled	0	0	none	
IsisL2	Disabled	0	0	none	
IsisL1Ext	Disabled	0	0	none	
IsisL2Ext	Disabled	0	0	none	
bgp	Disabled	0	0	none	

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Edge, Advanced Edge, or Core license. For licensing information, see the [Switch Engine 32.3 Feature License Requirements](#) document.

show ripng interface

```
show ripng interface {detail | vlan vlan-name | tunnel tunnel-name}
```

Description

Displays *RIPng*-specific configuration and statistics for the specified interface.

Syntax Description

detail	Specifies detailed display.
<i>vlan-name</i>	Specifies an IPv6 configured <i>VLAN</i> .
<i>tunnel-name</i>	Specifies an IPv6 tunnel.

Default

Show summary output for all interfaces.

Usage Guidelines

Displays the RIPng interface configuration and runtime information. If no interface is specified, only the summary data for all the configured interfaces is displayed. If an interface is specified, only the data for that interface is displayed in detail. If the keyword detail is specified, detailed data for all interfaces is displayed.

Example

The following command displays the RIPng configuration summary for all interfaces:

```
show ripng interface
```

The following is sample output from this command:

VLAN	IP Address	Flags	Sent Packets	Rcvd Packets	Triggered Updates	Cost
v1	22cc::3	/64 rif-pst	106349	106349	3	15
v2	22bb::1	/64 rif-pst	106349	106095	3	1
v3	2abc::1	/120 rif-pst	106351	0	4	1
v4	3ffe::1	/64 rif-pst	106349	139124	3	1

Flags: (f) Interface Forwarding Enabled, (i) Interface RIPng Enabled
 (n) Multinetted Interface, (r) Router RIPng Enabled
 (p) Poison Reverse Enabled, (s) Split Horizon Enabled
 (t) Triggerred Update Enabled.

The following command displays RIPng-specific statistics for the VLAN v1:

```
show ripng interface v1
```

The following is sample output from this command:

VLAN	: v1	Interface	: 22cc::3/64
Router RIPng	: Enabled	Cost	: 15
Input Policy	: None	Output Policy	: None
Trusted GW Policy	: gw6	Poison Reverse	: Enabled
Split Horizon	: Enabled	Triggered Updates	: Enabled
Rcvd Packets	: 106358	Sent Packets	: 106358
Sent Trig. Updates	: 3	Rcvd Bad Packets	: 0
Rcvd Bad Routes	: 0		
Neighbor Addresses	: fe80::201:30ff:fe94:f400		
Interface Addresses	: 22cc::3/64, fe80::280:c8ff:feb9:2855/64		

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Edge, Advanced Edge, or Core license. For licensing information, see the [Switch Engine 32.3 Feature License Requirements](#) document.

show ripng routes

```
show ripng routes {detail} {network ripngNetworkPrefix}
```

Description

Displays all matching routes in the *RIPng* routing database.

Syntax Description

detail	Displays all available information from the RIPng routing table.
ipv6-prefix	Specifies the route prefix for the routes to show.
prefix-length	Specifies the address mask of the IPv6 prefix.

Default

N/A.

Usage Guidelines

The routes displayed include all routes advertised by RIPng, including routes exported from the system routing table and originated by other protocols, for example OSPFv3 (also called redistributed routes).

Example

The following command displays a summary of RIPng specific routes:

```
show ripng routes
```

The following is sample output from this command:

```

Network                Next Hop                Mtr  VLAN
*> 2aaa::/64            fe80::201:30ff:fef4:5ca0%v1  2    v1
*                        fe80::201:30ff:fe94:f400%v2  2    v2
*> 2bbb::/64            fe80::201:30ff:fef4:5ca0%v1  2    v1
*                        fe80::201:30ff:fe94:f400%v2  3    v2
*> 2ccc::/64            (local)                    1    (direct)
*                        fe80::201:30ff:fef4:5ca0%v1  2    v1
*                        fe80::201:30ff:fe94:f400%v2  3    v2
*> 2ddd::/64            (local)                    1    (direct)
*                        fe80::201:30ff:fe94:f400%v2  2    v2

```

The following command displays the detailed RIPng route information:

```
show ripng routes detail
```

The following is sample output from this command:

```

IPv6 RIPng routing table entry for 2aaa::/64
Paths: (2 available, best #1)
    fe80::201:30ff:fef4:5ca0%v1 from fe80::201:30ff:fef4:5ca0%v1 (v1)
        Metric 2, tag 0, timeout in 02:38, valid, best
    fe80::201:30ff:fe94:f400%v2 from fe80::201:30ff:fe94:f400%v2 (v2)
        Metric 2, tag 0, timeout in 02:44, valid
IPv6 RIPng routing table entry for 2bbb::/64
Paths: (2 available, best #1)
    fe80::201:30ff:fef4:5ca0%v1 from fe80::201:30ff:fef4:5ca0%v1 (v1)
        Metric 2, tag 0, timeout in 02:38, valid, best
    fe80::201:30ff:fe94:f400%v2 from fe80::201:30ff:fe94:f400%v2 (v2)
        Metric 3, tag 0, timeout in 02:44, valid
IPv6 RIPng routing table entry for 2ccc::/64

```

```
Paths: (3 available, best #1)
  Local from direct
    Metric 1, tag 0, no timeout, valid, best
    fe80::201:30ff:fe4:5ca0%v1 from fe80::201:30ff:fe4:5ca0%v1 (v1)
    Metric 2, tag 0, timeout in 02:38, valid
    fe80::201:30ff:fe94:f400%v2 from fe80::201:30ff:fe94:f400%v2 (v2)
    Metric 3, tag 0, timeout in 02:44, valid
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Edge, Advanced Edge, or Core license. For licensing information, see the [Switch Engine 32.3 Feature License Requirements](#) document.

show rmon memory

```
show rmon memory {detail | memoryType}
```

Description

Displays RMON specific memory usage and statistics.

Syntax Description

detail	Displays detailed information.
<i>memoryType</i>	Specifies the type of memory usage and statistics to display.

Default

N/A.

Usage Guidelines

If you do not specify the detailed keyword or enter a specific RMON memory type, the output contains usage information for all memory types.

Example

The following command displays RMON memory statistics:

```
show rmon memory
```

The following is sample output from this command:

```
RMON Memory Information
-----
```

```

Bytes Allocated: 14298032 AllocFailed: 0
Current Memory Utilization Level: GREEN
Memory Utilization Statistics
-----
Size      16      32      48      64      80      96      112      128      144      176      208
256      384      5
12       768     1024     2048     4096     8192     16384     18432     40960     64000
-----
-----
Used Blocks 1558      3 2490      1      0      0      0      1      1      0
63444      1 1869
0      311      0      0      0      0      0      0      0      0
rmonEstat      0      0      0      0      0      0      0      0      0      0
0      0      311
0      0      0      0      0      0      0      0      0
rmonOwner 1555      0      0      0      0      0      0      0      0      0
0      0      0
0      0      0      0      0      0      0      0      0
rmonHisc      0      0      0      0      0      0      0      0      0
0      0      1244
0      0      0      0      0      0      0      0      0
rmonHist      0      0      0      0      0      0      0      0      0
63444      0      0
0      0      0      0      0      0      0      0      0
rmonAlarm      0      0      0      0      0      0      0      0      0
0      0      3
0      0      0      0      0      0      0      0      0
rmonLogDescription      0      0      0      0      0      0      0      0      1
0      0      0      0
0      0      0      0      0      0      0      0
rmonLog      0      1      0      0      0      0      0      0      0      0
0      0
0      0      0      0      0      0      0      0
rmonEvent      0      0      0      0      0      0      0      1      0      0
0      0      0
0      0      0      0      0      0      0      0
rmonEventDescription      0      1      0      0      0      0      0      0      0
0      0      0
0      0      0      0      0      0      0      0
rmonEventCommunity      0      1      0      0      0      0      0      0      0
0      0      0      0
0      0      0      0      0      0      0
rmonCommunity      1      0      0      0      0      0      0      0      0
0      0      0
0      0      0      0      0      0      0
rmonDs      0      0      0      0      0      0      0      0      0
0      0
0      311      0      0      0      0      0      0      0
rmonDbx      0      0 2490      0      0      0      0      0      0
0      0
0      0      0      0      0      0      0
rmonOid      0      0      0      0      0      0      0      0      0
0      311
0      0      0      0      0      0      0
rmonMdbIndexOid      2      0      0      1      0      0      0
0      0      0
0      0      0      0      0      0      0
rmonMdbString      0      0      0      0      0      0      0      0      0
0      1      0
0      0      0      0      0      0      0

```

The following command displays RMON event statistics:

```
show rmon memory rmonEvent
```

The following is sample output from this command:

```

RMON Memory Information
-----
Bytes Allocated: 14298032 AllocFailed: 0
Current Memory Utilization Level: GREEN
Memory Utilization Statistics
-----
Memory Statistics for rmonEvent
-----
Size      16      32      48      64      80      96      112      128      144      176      208
256      384      512      768      1024     2048     4096     8192     16384     18432     40960     64000
-----
-----
Allocated    0      0      0      0      0      0      0      0      1      0      0      0
0      0
0      0      0      0      0      0      0      0      0      0      0
AllocatedPeak    0      0      0      0      0      0      0      0      1      0      0
0      0      0
0      0      0      0      0      0      0      0      0      1      0      0
AllocSuccess    0      0      0      0      0      0      0      0      0      0      0
0      0      0      0      0      0      0      0      0      0
FreeSuccess    0      0      0      0      0      0      0      0      0      0      0
0      0      0
0      0      0      0      0      0      0      0      0      0
AllocFail    0      0      0      0      0      0      0      0      0      0      0
0      0      0
0      0      0      0      0      0      0      0      0
FreeFail    0      0      0      0      0      0      0      0      0      0      0
0      0      0
0      0      0      0      0      0      0      0      0
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show router-discovery

```
show router-discovery {ipv6} {vlan vlan_name}
```

Description

Displays the router discovery settings.

Syntax Description

<code>vlan_name</code>	Specifies an IPv6 configured VLAN .
------------------------	---

Default

N/A.

Usage Guidelines

If no VLAN is specified, the settings are displayed for all IPv6 configured VLANs.

Example

The following example displays router discovery settings for the VLAN "top_floor":

```
# show router-discovery vlan top_floor
Router Advertisements disabled on vl
Router Advertisements are sent only with VRRP's Virtual LLA: On
Minimum/Maximum Interval: 200 / 600
Managed / Other Info Flags: Off / Off
Link MTU: 0
Reachable Time: 0
Retrans Timer: 0
Current Hop Limit: 64
Default Lifetime: 1800
Number of Prefixes: 1, Prefix List:
Valid Preferred
Prefix Lifetime Auto Lifetime OnLink
2001:db8::1/64
2592000 On 604800 On
```

History

This command was first available in ExtremeXOS 11.2.

Router advertisements sent only with VRRP's virtual LLA information added in ExtremeXOS 32.3.

Platform Availability

This command is available on all platforms that support the Advanced Edge License as shown in the [Switch Engine 32.3 Feature License Requirements](#) document.

show rtep l2pt

```
show [[{vlan} vlan_name vman_name] [{vxlan {vr vr_name} rtep
    rtep_ip4}}] l2pt {detail}
```

Description

Displays the RTEP L2PT configuration and status of a service.

Syntax Description

vlan	Specifies the <u>VLAN</u> configuration.
<i>vlan_name</i>	Specifies the VLAN name.
vxlan	Specifies Virtual eXtensible LAN.
vr	Specifies Virtual Router.
<i>vr_name</i>	Specifies the Virtual Router Name. If not specified, the VR of the current command context is used.
rtep	Specifies Remote Tunnel End Point.
<i>rtep_ipv4</i>	Specifies the Remote Tunnel End Point IPv4 address.
l2pt	Specifies Layer 2 protocol tunneling.
detail	Specifies to show L2PT configuration and status in detail.

Default

N/A.

Usage Guidelines

Use this command to display the RTEP L2PT configuration and status of a service.

Example

The following example illustrates the `show rtep l2pt` command with RTEP IP address as 2.2.2.2 of VxLAN service Interface:

```
# show vlan tenant vxlan rtep 2.2.2.2 l2pt
Interface                                     Profile Name
-----
VR-Default:2.2.2.2                          l2pt-none
```

The following is an example of the `show rtep l2pt` command:

```
# show vlan tenant l2pt
Interface                                     Profile Name
-----
15                                           l2pt-user
VR-Default:2.2.2.2                          l2pt-none
```

The following example illustrates the `show rept l2pt detail` command with RTEP IP address as 2.2.2.2 of VxLAN service Interface:

```
# show vlan tenant vxlan rtep 2.2.2.2 l2pt detail
VXLAN RTEP: VR-Default:2.2.2.2
  L2PT Profile Name      : l2pt-none

  Protocol Filter Name   : cdp
  Destination Address:   01:00:0c:cc:cc:cc
  Protocol Id Type       : snap
  Protocol Id Value      : 0x2000
  Field Offset           :
  Field Value            :
```

```

Field Mask      :
Action          : None
CoS             :
DSCP            : 50
DSCP Replace    : Yes
Packets Transmitted: 0
Packets Received  : 0

```

The following example illustrates the `show rtep l2pt detail` command:

```

# show vlan tenant l2pt detail
Port 15
  L2PT Profile Name      : l2pt-user

  Protocol Filter Name   : cdp
  Destination Address    : 01:00:0c:cc:cc:cc
  Protocol Id Type       : snap
  Protocol Id Value      : 0x2000
  Field Offset           :
  Field Value            :
  Field Mask             :
  Action                 : Tunnel
  CoS                    :
  DSCP                   : 50
  DSCP Replace           : Yes
  Packets Transmitted    : 10956
  Packets Received       : 11492

VXLAN RTEP: VR-Default:2.2.2.2
  L2PT Profile Name      : l2pt-none

  Protocol Filter Name   : cdp
  Destination Address    : 01:00:0c:cc:cc:cc
  Protocol Id Type       : snap
  Protocol Id Value      : 0x2000
  Field Offset           :
  Field Value            :
  Field Mask             :
  Action                 : None
  CoS                    :
  DSCP                   : 50
  DSCP Replace           : Yes
  Packets Transmitted    : 0
  Packets Received       : 0

```

History

This command was first available in ExtremeXOS 31.3.

Platform Availability

This command is supported on the ExtremeSwitching 5520 series switch and stacks with 5520 slots only.

show script output autoexec

```
show script output autoexec
```

Description

Shows the results of executing the autoexec script.

Syntax Description

This command has no arguments or variables.

Default

N/A

Usage Guidelines

Use this command to show results when a autoexec.xsf file is executed. The file is not executed when a default.xsf file has been executed.

The CLI script file autoexec.xsf is executed after the configuration has been loaded. Its purpose is to run some commands after every reboot. It can also be used to revert to the original configuration following changes made by UPM executed persistent commands.

Example

This command shows the results of executing the autoexec script:

```
show script output autoexec
```

When there is no autoexec.xsf file, there is no response.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show script output default

```
show script output default
```

Description

Shows the results of executing default.xsf on bootup.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to show results when a default.xsf file is loaded.

An existing default.xsf file is executed if the switch comes up in an unconfigured state because the configuration file is missing, or the configuration file cannot be determined due to a corrupt NVRAM or other problems. This returns the switch to some basic configuration. When default.xsf is executed, the show switch command shows default.xsf as the booted configuration file.

Example

This command shows the results of executing the autoexec script:

```
show script output default
```

When there is no default.xsf file, there is no response.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show security

```
show security [fips-mode | python | tpm]
```

Description

Use this command to show FIPS mode, Trusted Platform Module (TPM), and external Python scripting support status.

Syntax Description

fips-mode	Shows specifically FIPS mode status.
python	Shows specifically external Python scripting support status.
tpm	Shows specifically X.509 certificates sorted in the switch's TPM chip.

Default

N/A

Usage Guidelines

If you select neither keyword option (FIPS/Python), you see status information for both.

For both FIPS mode or Python, two values appear:

- **Current**—shows the current active setting.
- **Configured**—show the setting that only takes effect after reboot.

If you select keyword **tpm** and the available *certificate* option, you will be presented with X.509 certificates that are provisioned in the TPM hardware: the Endorsement Key (EK), the Initial Attestation Key (IAK) and the Initial Device Identifier (IDevID) certificates. The EK is provisioned and signed by the TPM manufacturer, and the IAK and IDevID are provisioned and signed by Extreme Networks.

If the *text* option is not specified, then the certificate's PEM data will be displayed. If the *text* option is specified, then a human readable version of the certificate will be displayed.



Note

These certificates are informational only and currently not used.

Example

The following example shows both FIPS and Python scripting status:

```
# show security
FIPS Mode (current)      : Off
FIPS Mode (configured)  : On
Python (current)         : Off
Python (configured)     : On
```

The following example shows only Python scripting status:

```
# show security python
Python (current)         : Off
Python (configured)     : Off
```

The following example shows the TPM certificate options:

```
# show security tpm certificate
ek          Endorsement Key certificate
iak         Initial Attestation Key certificate
idevid      Initial Device Identifier certificate
```

The following is an example EK certificate with both RSA and ECC keys:

```
# show security tpm certificate ek

[Endorsement Key RSA Certificate]

-----BEGIN CERTIFICATE-----
MIIEJzCCA3egAwIBAgIEJfR5OjANBgkqhkiG9w0BAQsFADB2MQswCQYDVQQGEwJE
RTEhMB8GA1UECgwYSW5maW5lb24gVG9jaG5vbG9naWVzIEFHRMRwEQYDVQQLDAP
UFRJR0EoVE0pMS8wLQYDVQQDDCZJbmZpbmVvbiBPUFRJR0EoVE0pIFRQTSAYLjAg
U1NBIE5BIDA0MjAeFw0xOTA5MDMwODM5MTNaFw0zNDA5MDMwODM5MTNaMAAwggEi
MA0GCSqGSIb3DQEBBQUAA4IBDwAwggEKAoIBAQC0mU058b7ZYzkqAsldjeoYVeFU
GBQi1Ce0x3bjAQql7SW6YbeJtZZ8mj8mWENoUO7X3lMBNABf040IOBRq+fhqaSlK
M0UAiWzgzSZjqFuQGLdZsv3aK1g89eKFaBRzZUfqH4bHcqzU6dNBQ+Zj6IpTd0Kus
x11zPqCwBEO39SnUtlalOee72qaHeXqS6GPioNseJSnHOYMI36zM4JVzJ4rTr17J
```

```
f+jslVnpaghCOzNlbypl7DyzgnbWn2slLxvoN/+0+0g1kzEclPoS/nfft2q4Gnov
pmGwPHN8i0OerCKwXkuo7pAjvz9Q0++06igMiAeEHhui62CxU8bkdeFdINf5AgMB
AAGjggGZMIIBlTBbBggrBgEFBQcBAQRPMEOwSwYIKwYBBQUHMAKGP2h0dHA6Ly9w
a2kuaW5maW51b24uY29tL09wdGlnYVJzYU1mckNBMDQyL09wdGlnYVJzYU1mckNB
MDQyLmNyddAOBGNVHQ8BAf8EBAMCACAwwAYDVR0RAQH/BE4wTKRKMegxFjAUBgVn
gQUCAQwLawQ6NDk0NjU4MDAxGjAYBgVngQUCAgwPU0xCIDk2NzAgVFBnMi4wMRIw
EAYFZ4EFAGMMB21kOjA3NTUwDAYDVR0TAQH/BAIwADBQBgNVHR8ESTBHMEWgQ6BB
hj9odHRwOi8vcGtpLmluZmluZW9uLmNvbS9PcHRpZ2FSc2FNZnJDQTA0Mi9PcHRp
Z2FSc2FNZnJDQTA0Mi5jcmwwFQYDVR0gBA4wDDAKBggqghQARAEUATAfBgNVHSME
GDAWgBRdCBWVH19gY4pp5yUvPsS+zXVUsjAQBgNVHSUECTAHBgVngQUIATAiBgNV
HQQEGZAZMBcGBWeBBQIQM4wDAwDmi4wAgEAAgIAIjANBgkqhkiG9w0BAQsFAAOC
AQEAT7fhElXcMITmsF6pC2xtryszIu2Gq7618+fDoiOI8Qvvc2pD4BK5i+U+tjCW
UwfJxB9v86tSs9Fvh2PWmC36k58+Gkz/04yBlr15vLcgnkEr38dFKr4PkQULkbiK
t1FATPMbbj9NY4xJlLxwOcTsrzn0EkCqLUiVDUH3ohMppjQMPiKL/zS/t/aiAUsoQ
8po3cNkuPv/hUgKzhPPetKUpIVz1NLatmz052N5kqabjd4EwDLkXrDVoOIR8SRWa
8xHBGBxTkWqAgv/UVgl6kDF0JsteDvH//oU5+MbAx9PWQv3cddQgcZiKeO1qNHMB
0Tj4FREumRw7L1lQb3/hUKIH0Q==
-----END CERTIFICATE-----
```

```
[Endorsement Key ECC Certificate]
```

```
-----BEGIN CERTIFICATE-----
MIIDBDCAQmgAwIBAgIERslzNTAKBggqhkjOPQQDAjB2MQswCQYDVQGEwJERTEh
MB8GA1UECgwYISW5maW51b24gVGVjaG5vbG9naWVzIEFHMHRMwEQYDVQLDAPUFRJ
R0EOvE0pMS8wLQYDVQDDCZJbmZpbmVvbiBPUFRJR0EOvE0pIFRQTSAYLjAgRUND
IENBIDA0MjAeFw0xOTA5MDMwODM4NTNaFw0zNDA5MDMwODM4NTNaMAAwWTATBgCq
hkjOPQIBBggqhkjOPQMBBwNCAAAQC24h7AgT0ZL/wOgwN+47R8NjHtddRnyroalsk
/x/m4mLdEXqGD7913Tt/d9QhGafOoUkgIsOVLNfw4fy0ZNPvo4IBmTCCAZUwWwYI
KwYBBQUHAQEETzBNMESGCCSGAUQFBzAchj9odHRwOi8vcGtpLmluZmluZW9uLmNv
bs9PcHRpZ2FFY2NNZnJDQTA0Mi9PcHRpZ2FFY2NNZnJDQTA0Mi5jcnQwDgYDVR0P
AQH/BAQDAGAIMFgGA1UdEQEB/wROMEYksjBIMRYwFAYFZ4EFAGEMC21kOjQ5NDY1
ODAwMR0wGAYFZ4EFAGIMD1NMQIA5NjcwIFRQTTIuMDESMBAGBWeBBQIDDApZDow
NzU1MAwGA1UdEwEB/wQMAAwUAYDVR0fBEKwRzBfEOgQYY/aHR0cDovL3BraS5p
bmZpbmVvbi5jb20vT3B0aWdhRWVjZmZlZmZlZmZlZmZlZmZlZmZlZmZlZmZlZmZl
Y3JsbmZpbmVvbi5jb20vT3B0aWdhRWVjZmZlZmZlZmZlZmZlZmZlZmZlZmZlZmZl
nCc6QWjtUP1JEl5EwEAYDVR0lBAkwBwYFZ4EFCAEwIgYDVR0lJBBSwGTAXBgVngQUC
EDEOMAwMAzIuMAIBAAICAIAowCgYIKoZIzj0EAwIDSQAARgIhAL//3+inIwQg/gOh
cWotTy2FaQ8NdpYDi4LYPtFwIpXpAiEAx2m6Q4oIvf0EIwKqzD684kkezcoubm/
KbaiaGUA4x8=
-----END CERTIFICATE-----
```

The following is an example EK certificate with the human-readable *text* option with both RSA and ECC keys:

```
# show security tpm certificate ek text
```

```
[Endorsement Key RSA Certificate]
```

```
Certificate:
```

```
Data:
```

```
Version: 3 (0x2)
```

```
Serial Number: 636778810 (0x25f4793a)
```

```
Signature Algorithm: sha256WithRSAEncryption
```

```
Issuer: C=DE, O=Infineon Technologies AG, OU=OPTIGA(TM), CN=Infineon OPTIGA(TM)
```

```
TPM 2.0 RSA CA 042
```

```
Validity
```

```
Not Before: Sep 3 08:39:13 2019 GMT
```

```
Not After : Sep 3 08:39:13 2034 GMT
```

```
Subject:
```

```
Subject Public Key Info:
```

```
Public Key Algorithm: rsaEncryption
```

```
Public-Key: (2048 bit)
```

```
Modulus:
```

```

00:b4:99:4d:39:f1:be:d9:61:99:20:02:cd:5d:8d:
ea:18:55:e1:54:18:14:22:d4:27:b4:c7:76:e3:01:
0a:a5:ed:25:ba:61:b7:89:b5:96:7c:9a:3f:26:58:
43:68:50:ee:d7:df:59:81:34:00:5f:d3:8d:08:38:
14:6a:f9:f8:6a:69:29:4a:33:45:00:8b:0c:e0:49:
98:ea:16:e4:06:2d:d6:6c:bf:76:8a:d6:0f:3d:78:
a1:5a:05:1c:d9:51:fa:87:e1:b1:dc:ab:35:3a:74:
d0:50:f9:98:fa:22:94:dd:d0:ab:92:c6:5d:73:3e:
a0:b0:04:43:b7:f5:29:d4:b6:36:a5:39:e7:bb:da:
a6:87:79:7a:92:e8:63:e2:a0:db:1e:25:29:c7:39:
83:08:df:ac:cc:e0:95:73:27:8a:d3:af:5e:c9:7f:
e8:ec:d5:59:e9:6a:08:42:3b:33:75:6f:2a:4b:ec:
3c:b3:82:76:d6:9f:6b:25:2f:1b:e8:37:ff:b4:fb:
48:35:93:31:1c:94:fa:12:fe:77:df:4f:6a:b8:1a:
7a:2f:a6:61:b0:3c:73:7c:8b:43:9e:ac:22:b0:5e:
4b:a8:ee:90:23:bd:9f:50:d3:ef:b4:ea:28:0c:88:
07:84:1e:1b:a2:eb:60:b1:53:c6:e4:74:47:dd:20:
d7:f9
Exponent: 65537 (0x10001)
X509v3 extensions:
  Authority Information Access:
    CA Issuers - URI:http://pki.infineon.com/OptigaRsaMfrCA042/
OptigaRsaMfrCA042.crt

  X509v3 Key Usage: critical
    Key Encipherment
  X509v3 Subject Alternative Name: critical
    DirName:/2.23.133.2.1=id:49465800/2.23.133.2.2=SLB 9670
TPM2.0/2.23.133.2.3=id:0755
  X509v3 Basic Constraints: critical
    CA:FALSE
  X509v3 CRL Distribution Points:

    Full Name:
      URI:http://pki.infineon.com/OptigaRsaMfrCA042/OptigaRsaMfrCA042.crl

  X509v3 Certificate Policies:
    Policy: 1.2.276.0.68.1.20.1

  X509v3 Authority Key Identifier:
    keyid:5D:08:15:95:1F:5F:60:63:8A:69:E7:25:2F:3E:C4:BE:CD:75:54:B2

  X509v3 Extended Key Usage:
    2.23.133.8.1
  X509v3 Subject Directory Attributes:
    0.0...g....1.0...2.0.....
Signature Algorithm: sha256WithRSAEncryption
4f:b7:e1:12:55:dc:30:84:e6:b0:5e:a9:0b:6c:6d:af:2b:33:
22:ed:86:ab:be:a5:f3:e7:c3:a2:23:88:9b:c4:2f:bd:cd:a9:
0f:80:4a:e6:2f:94:b6:30:96:53:07:c9:c4:1f:6f:f3:ab:52:
b3:d1:6f:87:63:d6:98:2d:fa:93:9f:3e:1a:4c:ff:d3:8c:81:
96:bd:79:bc:b7:20:9e:41:2b:df:c7:45:2a:be:0f:91:05:0b:
91:b8:8a:b7:51:40:4c:f3:1b:6e:3f:4d:63:8c:49:94:bc:70:
39:c4:ec:af:39:f4:12:40:aa:2d:48:95:0d:41:f7:a2:13:29:
8d:03:29:20:a2:ff:cd:2f:ed:fd:a8:80:52:c3:90:f2:9a:37:
70:d9:2e:3e:ff:e1:52:02:b3:84:f3:c4:b4:a5:29:21:5c:e5:
34:b6:ad:9b:3d:39:d8:de:64:a9:a6:e3:77:81:30:0c:b9:17:
ac:35:68:38:84:7c:49:15:9a:f3:11:c1:18:1c:53:93:0a:80:
82:ff:d4:56:09:7a:90:31:74:26:cb:5e:0e:f1:ff:fe:85:39:
f8:c6:c0:c7:d3:d6:42:fd:dc:75:d4:20:71:98:8a:78:ed:6a:
34:73:1b:d1:38:f8:15:11:2e:99:1c:3b:2e:5d:50:6f:7f:e1:
52:42:07:d1

```

```
[Endorsement Key ECC Certificate]

Certificate:
  Data:
    Version: 3 (0x2)
    Serial Number: 1187607349 (0x46c97335)
    Signature Algorithm: ecdsa-with-SHA256
    Issuer: C=DE, O=Infineon Technologies AG, OU=OPTIGA(TM), CN=Infineon OPTIGA(TM)
TPM 2.0 ECC CA 042
    Validity
      Not Before: Sep  3 08:38:53 2019 GMT
      Not After : Sep  3 08:38:53 2034 GMT
    Subject:
      Subject Public Key Info:
        Public Key Algorithm: id-ecPublicKey
        Public-Key: (256 bit)
        pub:
          04:02:db:88:7b:02:04:f4:64:bf:f0:3a:0c:0d:fb:
          8e:d1:f0:d2:61:4d:d7:51:9f:2a:e8:6b:5b:24:ff:
          1f:e6:e2:62:dd:11:7a:86:0f:bf:75:dd:3b:7f:77:
          d4:21:18:07:ce:a1:49:20:22:c3:95:2c:d7:f0:e1:
          fc:b4:64:d3:ef
        ASN1 OID: prime256v1
        NIST CURVE: P-256
    X509v3 extensions:
      Authority Information Access:
        CA Issuers - URI:http://pki.infineon.com/OptigaEccMfrCA042/
OptigaEccMfrCA042.crt

      X509v3 Key Usage: critical
      Key Agreement
      X509v3 Subject Alternative Name: critical
      DirName:/2.23.133.2.1=id:49465800/2.23.133.2.2=SLB 9670
TPM2.0/2.23.133.2.3=id:0755
      X509v3 Basic Constraints: critical
      CA:FALSE
      X509v3 CRL Distribution Points:

      Full Name:
      URI:http://pki.infineon.com/OptigaEccMfrCA042/OptigaEccMfrCA042.crl

      X509v3 Certificate Policies:
      Policy: 1.2.276.0.68.1.20.1

      X509v3 Authority Key Identifier:
      keyid:B1:1F:33:CC:A6:06:56:BA:25:9C:2E:90:5A:3B:54:3F:52:44:97:91

      X509v3 Extended Key Usage:
      2.23.133.8.1
      X509v3 Subject Directory Attributes:
      0.0...g....1.0...2.0.....
    Signature Algorithm: ecdsa-with-SHA256
      30:46:02:21:00:bf:ff:df:e8:a7:23:04:20:fe:03:a1:71:6a:
      2d:4f:2d:85:69:0f:0d:76:96:03:8b:82:d8:3e:d1:70:22:95:
      e9:02:21:00:c7:69:ba:43:8a:08:bd:fd:04:23:02:aa:cc:3e:
      bc:e2:49:1e:cd:ca:2e:6e:b9:bf:29:b6:a2:6a:05:00:e3:1f
```

History

This command was first available in ExtremeXOS 21.1.

External Python scripting support status was added in ExtremeXOS 32.3.

The **tpm** option was added in ExtremeXOS 31.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show session

```
show session [{detail} {sessID}] {history}
```

Description

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

Syntax Description

detail	Specifies more detailed session information.
<i>sessID</i>	Specifies a session ID number.
history	Displays a list of all sessions.

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. Each session is numbered.

Beginning with ExtremeXOS 11.2, the switch accepts IPv6 connections. If the incoming session is from an IPv6 address, the show session output indicates IPv6.

You can specify the following options to alter the session output:

- **detail**—The output for all current sessions is displayed in a list format.
- **sessID**—The output for the specified session is displayed in a list format.
- **history**—Displays a list of current and previous sessions, including the user, type of session, location, and start and end time of the session.

The `show session` command fields are defined in the following table.

Table 44: 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.

Table 44: Show Command Field Definitions (continued)

Field	Definition
Type	Indicates the type of session, for example: console, Telnet, HTTP, HTTPS.
Auth	Indicates how the user is logged in (local, <u>RADIUS</u> , TACACS+, sshKey, x509v3).
CLI Auth	Indicates the type of authentication (RADIUS and TACACS) if enabled.
Location	Indicates the location (IP address) from which the user logged in. The output also indicates if the location is an IPv6 address.

Example

The following command displays the active sessions on the switch:

```
show session
```

The following is sample output from this command:

```
CLI
#      Login Time          User      Type    Auth    Auth Location
=====
1      Thu Apr 28 20:16:56 2005 admin     console local   dis    serial
*2      Thu Apr 28 23:36:20 2005 admin     ssh2    local   dis    3001::20d:88ff:fec5:ad40
3      Fri Apr 29 11:14:27 2005 admin     telnet  local   dis    10.255.44.55
```

The following command displays a list of current and previous sessions on the switch:

```
show session history
```

The following is sample output from this command:

```
Session History:
admin                               console    serial
Mon Jun 21 09:19:00 2004
Mon Jun 21 10:00:16 2004
admin                               console    serial
Tue Jun 22 07:28:
11 2004
Tue Jun 22 11:46:48 2004
admin                               console    serial
Wed Jun 23 10:05:44 2004
Wed Jun 23 14:11:47 2004
admin                               console    serial
Thu Jun 24 07:07:25 2004
Thu Jun 24 07:08:55 2004
admin                               console    serial
Thu Jun 24 13:30:07 2004 Active
```

The following displays X509v3 Authentication information (lines 4 and 20):

```
# show session

#      Login Time          User      Type    Auth    CLI
#      Login Time          User      Type    Auth    Auth Location
=====
*1      Wed Dec 15 14:53:43 2021 vasanth  console RADIUS      dis    serial
4      Wed Dec 15 14:57:28 2021 exos-  .. ssh2    x509v3-local dis    10.120.89.84
* (pacman debug) 460-11_switch.9 #
```

```
* (pacman debug) 460-11_switch.9 # show session detail
Session Detail:
ID                :1
Start Time        :Wed Dec 15 14:53:43 2021
User Name         :vasanth
Type              :console
Authentication     :RADIUS
Cli Authentication :disabled
Location          :serial
Node              :local

ID                :4
Start Time        :Wed Dec 15 14:57:28 2021
User Name         :exos-admin
Type              :ssh2
Authentication     :x509v3-local
Cli Authentication :disabled
Location          :10.120.89.84
Node              :local
```

History

This command was first available in ExtremeXOS 10.1.

Support for IPv6 was added in ExtremeXOS 11.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show sflow configuration

```
show sflow {configuration}
```

Description

Displays the current sFlow configuration.

Syntax Description

configuration	Displays the sFlow configuration.
----------------------	-----------------------------------

Default

N/A.

Usage Guidelines

This command displays the sFlow configuration of your system.

The following fields are displayed:

- Global Status—sFlow is globally enabled or disabled.
- Polling interval—How often the hardware is polled for statistics, in seconds.
- Sampling rate—Packets are sampled, on average, once for every rate-number of packets.
- Maximum cpu sample limit—Maximum number of packets per second sampled before sample throttling takes effect.
- Agent IP—IP address inserted into the sFlow data packets to identify the sFlow switch.
- Collectors—To which IP address and port, and from which virtual router, the sFlow packets are sent.
- Port Status—Enabled or disabled for statistics gathering.
- Port Sample-rate—Shows the sampling rate configured for the port and the actual rate if CPU throttling has taken effect.
- Port Subsampling factor—for details, see the command [configure sflow max-cpu-sample-limit](#) on page 1234.

Example

To display the sFlow configuration on your system, use the following command:

```
show sflow
```

The following is an example of the show sflow configuration command :

```
SFLOW Global Configuration
Global Status: enabled
Polling interval: 20
Sampling rate: 8192
Maximum cpu sample limit: 2000
SFLOW Configured Agent IP: 0.0.0.0
Operational Agent IP: 10.127.11.88
Collectors

SFLOW Port Configuration
Port  Status          Sample-rate          Subsampling          Sflow-type
      Config / Actual      factor              Ingress / Egress
5:21  enabled          8192 / 8192         1                    Disabled / Enabled
```

History

This command was first available in an ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show sflow hardware-utilization

```
show sflow hardware-utilization {slot [slot_num | all]}
```

Description

Displays sFlow hardware table utilization statistics.

Syntax Description

hardware-utilization	Specifies showing hardware table utilization statistics.
slot	Specifies selecting slots.
<i>slot_num</i>	Specifies which slots to show statistics for.
all	Specifies showing statistics for all slots.

Default

N/A.

Usage Guidelines

This command displays sFlow hardware table utilization statistics for your system.

Example

The following example shows sFlow hardware table utilization statistics.

```
# show sflow hardware-utilization
sFlow Hardware Table Utilization Statistics
```

Resource Type	Current	Maximum	% Util.
MAC Entries	3	73727	0
Host Entries	4	73728	0
IPv4 Entries	0	73728	0
IPv6 Entries	2	36864	0
Long IPv6 Entries	0	2048	0
Total Routes	11	8192	0
IPv4 Neighbors	5	N/A	N/A
IPv6 Neighbors	0	N/A	N/A
IPv4 routes	3	N/A	N/A
IPv6 routes	3	N/A	N/A
ECMP Next Hops	2	1023	0
ACL Ingress Entries	42	10752	0
ACL Ingress Counters	0	233472	0
ACL Ingress Meters	0	12288	0
ACL Ingress Slices	0	12	0
ACL Egress Slices	0	4	0

```
Legend: N/A - Maximum not defined in published sFlow structure.
        Also see IP usage in 'show iproute reserved-entries statistics'.
```

History

This command was first available in ExtremeXOS 22.5.

Ability to select slots in a stack was added in ExtremeXOS 22.6.

Legend explaining N/A entries was added in ExtremeXOS 30.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show sflow statistics

```
show sflow statistics
```

Description

Displays sFlow statistics.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command displays sFlow statistics for your system.

The following fields are displayed:

- Received frames—Number of frames received on sFlow enabled ports.
- Sampled Frames—Number of packets that have been sampled by sFlow.
- Transmitted Frames—Number of UDP packets sent to remote collector(s).
- Broadcast Frames—Number of broadcast frames received on sFlow enabled ports.
- Multicast Frames—Number of multicast frames received on sFlow enabled ports.
- Packet Drops—Number of samples dropped.

Example

To display sFlow statistics for your system, use the following command:

```
show sflow statistics
```

The output from this command is similar to the following:

```
SFLOW Statistics
Received frames      : 1159044921
Sampled Frames       : 104944
Transmitted Frames   : 10518
Broadcast Frames     : 0
```

```
Multicast Frames      : 1055652
Packet Drops          : 0
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show sharing

```
show {port port_number} sharing {detail}
```

Description

Displays general loading sharing information.

Syntax Description

port	View load sharing information for a specific port.
<i>port_number</i>	View load sharing information for a specific port.
detail	Show details for the specified ports in the load sharing group.

Default

N/A.

Example

The following example shows load sharing information for the switch:

```
# show sharing
show sharing
Load Sharing Monitor
Config      Current Agg      Min      Ld Share  Dist  Ld Share  Agg Link  Link Up
Master      Master  Control Active Algorithm Flags Group      Mbr State Transitions
=====
  1:91  1:91      Static      1    L3_L4      A    1:91      Y    A      1
=====
Link State: A-Active, D-Disabled, R-Ready, NP-Port not present, L-Loopback
Minimum Active: (<) Group is down. # active links less than configured minimum
Load Sharing Algorithm: (L2) Layer 2 address based, (L3) Layer 3 address based
                      (L3_L4) Layer 3 address and Layer 4 port based
                      (custom) User-selected address-based configuration
Custom Algorithm Configuration: ipv4 L3-and-L4, xor
Custom Hash Seed: Switch MAC address (0x969C3CE2)
Distribution Flags:
  A - All: Distribute to all members,
  L - Local Slot: Distribute to members local to ingress slot,
  P - Port Lists: Distribute to per-slot configurable subset of members,
```

```
R - Resilient Hashing enabled.  
Number of load sharing trunks: 1
```

History

This command was first available in ExtremeXOS 11.0.

Custom hash seed information was added in ExtremeXOS 30.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show sharing distribution port-based

```
show sharing distribution port port-based {ports port_list}
```

Description

Displays the load sharing distribution to member ports in the group specified by *port* for traffic received on ports in *port_list*.

Syntax Description

ports	Specifies the group of member ports to display for the load sharing.
<i>port_list</i>	Specifies the group of ports from which traffic is received.

Default

N/A.

Usage Guidelines

Use this command to display the load sharing distribution to member ports in the group specified by *port* for traffic received on ports in *port_list*. The selected member ports displayed are the results of the calculation using the keys for the ports in the *port_list*, and the list of aggregator ports for the load sharing group. This command serves as port-based load sharing calculator for convenience.

Example

The following output shows the egress member ports selected for distribution in a load sharing group with master port 5:1, and aggregator ports 5:1 and 6:1 for packets received on ports 1:1-8, with default keys as shown.

```
# show sharing port-based keys ports 1:1-8  
1:1:  0   1:2:  1   1:3:  2   1:4:  3   1:5:  4   1:6:  5   1:7:  6   1:8:  7  
# show sharing distribution 5:1 port-based keys ports 1:1-8  
1:1 ->  5:1  
1:2 ->  6:1
```

```

1:3 -> 5:1
1:4 -> 6:1
1:5 -> 5:1
1:6 -> 6:1
1:7 -> 5:1
1:8 -> 6:1

```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show sharing health-check

```
show sharing health-check
```

Description

Displays the configured health check LAGs on a switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to display the health-check LAGs that have been configured on the switch.

Example

The following is sample output from this command:

```

# show sharing health-check
Member  Agg Admin Track      Track
Group   Port  Mbr State IP Addr      TCP Port Miss Freq State  Dn  Up
=====
2:8      2:1*  Y   En   30.1.1.1      23          3   3   Up    0   1
2:2*  Y   En   30.1.1.2      23          3   3   Up    0   1
2:3*  Y   En   30.1.1.3      23          3   3   Up    0   1
2:8*  -   En   30.1.1.8      80          3  10  Down   0   0
2:11*  Y   -   -             -            -   -   -    -   -
2:12*  -   En   44.1.3.2      80          3   4  Down   0   0
2:16  -   En   30.1.1.16     80          3  10  Dis    0   0
2:20    2:20* Y   En   192.1.1.1     80          10  3   Up    0   1
2:21*  Y   En   192.1.1.2     80         10   3   Up    0   1

```

```
=====
Member Port Flags: (*)Active, (!) Disabled
```

History

This command was first available in ExtremeXOS 12.13.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show sharing port-based keys

```
show sharing port-based keys {ports port_list}
```

Description

Displays the load sharing key values for all ports in the *port_list*. These values may be either default values, or configured values.

Syntax Description

ports	Specifies the ports for the load sharing.
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

N/A.

Usage Guidelines

Use this command to display the health-check LAGs that have been configured on the switch.

Example

```
X460G2-24p-24hp.1 # show sharing port-based keys
 1: 0    2: 1    3: 2    4: 3    5: 4    6: 5    7: 6    8: 7
 9: 8   10: 9   11: 10   12: 11   13: 12   14: 13   15: 14   16: 15
17: 0   18: 1   19: 2   20: 3   21: 4   22: 5   23: 6   24: 7
25: 8   26: 9   27: 10   28: 11   29: 12   30: 13   31: 14   32: 15
33: 0   34: 1   35: 2   36: 3   37: 4   38: 5   39: 6   40: 7
41: 8   42: 9   43: 10   44: 11   45: 12   46: 13   47: 14   48: 15
49: 0   50: 1   51: 2   52: 3
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show slot

```
show slot {slot {detail} | detail }
```

Description

Displays the slot-specific information.

Syntax Description

<i>slot</i>	Specifies a slot on a SummitStack or bridge port extender (BPE) bridge port extender (BPE) slot assignment.
detail	Specifies detailed port information.

Default

N/A.

Usage Guidelines

The `show slot` command displays the following information:

- The slot number.
- The type of switch installed in the slot.
- The type of switch configured for the slot.
- The state of the switch, whether the power is down, if the switch is operational, if a diagnostic being run, if there is a mismatch between the slot configuration and the switch in the slot.
- The number of ports on the switch.
- The current number of times the switch has been restarted after a failure and the configured restart-limit.



Note

You may see slightly different information displayed depending on the platform and configuration you are using.

If you do not specify a slot number, information for all slots is displayed.

The display also includes a notice of insufficient power, should that arise.

The `show slot` command displays the following states, among others:

- Down
- Power ON
- Powered OFF
- Booting

- Initializing
- VLAN sync
- FDB sync
- ACL sync
- RT sync
- Operational

The output of this command displays eight rows. If a node in the Active Topology is not assigned a slot number, the state of the slot is shown as Empty. A node that shows as a slot has successfully joined the active topology. This means ExtremeXOS software can communicate with this node and does not mean that the node has been successfully brought up. The card state in the display indicates whether the slot was successfully started. If the card state is Operational, then the node is being used in the stack to carry user data as configured.

The output of the command displays the slot number, type of the ExtremeSwitching in that slot, state of the slot, and the number of ports.

The number of ports does not include the stacking links. It includes the option card ports regardless of whether the option card is installed.

This command is not available on ExtremeSwitching series switches operating in non-stacking mode.

Example

The following example displays switch information for all slots in a stack:

```
* Slot-7 Stack.1 # show slot
Slots      Type              Configured              State      Ports
-----
Slot-1     SummitX                  SummitX                  Operational 26
Slot-2     X                        X                        Operational 50
Slot-3     X                        X                        Operational 26
Slot-4     SummitX                  SummitX440               Operational 26
Slot-5     X                        X                        Operational 26
Slot-6     X                        X                        Operational 26
Slot-7     X                        X                        Operational 26
Slot-8                                           Empty          0
```

The following example displays switch information for a specific slot on the stack:

```
* Slot-7 Stack.91 # show slot 1 detail
Slot-1 information:
State:                Operational
Download %:           100
Restart count:        0 (limit 5)
Serial number:         800187-00-02 0635G-00074
Hw Module Type:       SummitX450-24x
SW Version:           12.0.0.17
SW Build:              v1170b17
Configured Type:       SummitX440
Ports available:       26
Recovery Mode:         Reset
Node MAC:              02:04:96:27:87:17
Current State:         STANDBY
Image Selected:        secondary
Image Booted:          secondary
```

```

Primary ver:      12.0.0.16
Secondary ver:    12.0.0.17
Config Selected:  primary.cfg

```

The following example displays detailed switch information for all slots on a stack:

```

* Slot-7 Stack.90 # show slot detail
Slot-1 information:
State:                Operational
Download %:           100
Restart count:        0 (limit 5)
Serial number:        800187-00-02 0635G-00074
Hw Module Type:       SummitX
SW Version:           12.0.0.17
SW Build:             v1170b17
Configured Type:      SummitX
Ports available:      26
Recovery Mode:        Reset
Node MAC:             02:04:96:27:87:17
Current State:        STANDBY
Image Selected:       secondary
Image Booted:         secondary
Primary ver:          12.0.0.16
Secondary ver:        12.0.0.17
Config Selected:      primary.cfg
Slot-2 information:
State:                Operational
Download %:           100
Restart count:        0 (limit 5)
Serial number:        800163-00-04 0635G-01187
Hw Module Type:       SummitX
SW Version:           12.0.0.17
SW Build:             v1170b17
Configured Type:      SummitX
Ports available:      50
Recovery Mode:        Reset
Node MAC:             02:04:96:27:87:17
Current State:        STANDBY
Image Selected:       secondary
Image Booted:         secondary
Primary ver:          12.0.0.16
Secondary ver:        12.0.0.17
Config Selected:      primary.cfg
Slot-3 information:
State:                Operational
Download %:           100
Restart count:        0 (limit 5)
Serial number:        800152-00-04 0630G-00736
Hw Module Type:       SummitX
SW Version:           12.0.0.17
SW Build:             v1170b17
Configured Type:      SummitX
Ports available:      26
Recovery Mode:        Reset
Node MAC:             02:04:96:27:87:17
Current State:        STANDBY
Image Selected:       secondary
Image Booted:         secondary
Primary ver:          12.0.0.16
Secondary ver:        12.0.0.17
Config Selected:      primary.cfg
Slot-4 information:
State:                Operational
Download %:           100

```

```
Restart count:      0 (limit 5)
Serial number:      0635G-00073 S450-24X
Hw Module Type:     SummitX
SW Version:          12.0.0.17
SW Build:            v1170b17
Configured Type:     SummitX
Ports available:     26
Recovery Mode:       Reset
Node MAC:            02:04:96:27:87:17
Current State:       STANDBY
Image Selected:      secondary
Image Booted:        secondary
Primary ver:         12.0.0.16
Secondary ver:       12.0.0.17
Config Selected:     primary.cfg
Slot-5 information:
State:               Operational
Download %:          100
Restart count:       0 (limit 5)
Serial number:       800153-00-04 0646G-00683
Hw Module Type:      SummitX
SW Version:          12.0.0.17
SW Build:            v1170b17
Configured Type:     SummitX
Ports available:     26
Recovery Mode:       Reset
Node MAC:            02:04:96:27:87:17
Current State:       STANDBY
Image Selected:      secondary
Image Booted:        secondary
Primary ver:         12.0.0.16
Secondary ver:       12.0.0.17
Config Selected:     primary.cfg
Slot-6 information:
State:               Operational
Download %:          100
Restart count:       0 (limit 5)
Serial number:       800153-00-04 0646G-00691
Hw Module Type:      SummitX
SW Version:          12.0.0.17
SW Build:            v1170b17
Configured Type:     SummitX
Ports available:     26
Recovery Mode:       Reset
Node MAC:            02:04:96:27:87:17
Current State:       BACKUP
Image Selected:      secondary
Image Booted:        secondary
Primary ver:         12.0.0.16
Secondary ver:       12.0.0.17
Config Selected:     primary.cfg
Slot-7 information:
State:               Operational
Download %:          100
Restart count:       0 (limit 5)
Serial number:       800153-00-01 0603G-00741
Hw Module Type:      SummitX
SW Version:          15.5.0.0
SW Build:            v1170b17
Configured Type:     SummitX
Ports available:     26
Recovery Mode:       Reset
Node MAC:            02:04:96:27:87:17
Current State:       MASTER
```

```

Image Selected:      secondary
Image Booted:        secondary
Primary ver:         12.0.0.16
Secondary ver:       12.0.0.17
Config Selected:     primary.cfg
Slot-8 information:
State:               Empty
Restart count:       0 (limit 5)
Serial number:
Hw Module Type:
Configured Type:
Ports available:     0
Recovery Mode:       Reset
Node MAC:            00:00:00:00:00:00
Current State:
Image Selected:
Image Booted:
Primary ver:
Secondary ver:
Config Selected:

```

The following example shows information for a BPE on slot 100:

```

# show slot 100 detail

Slot-100 information:
  State:               Operational
  Description:         Building1, Floor30
  Download %:          100
  Flags:               M
  Restart count:       0 (limit 5)
  Serial number:       00.00.01 1705D-10009
  Hw Module Type:      V400-48t-10GE4
  SW Version:          1.1.0.32
  Configured Type:     V400-48t-10GE4
  Ports available:     52
  Recovery Mode:       Reset
  Debug Data:          Peer=

```

History

This command was first available in ExtremeXOS 10.1.

This command was first available on SummitStack in ExtremeXOS 12.0.

Slot description (name) information was added in ExtremeXOS 22.5.

Platform Availability

This command is available only on SummitStacks and in VPEx mode.

show slpp guard

```
show slpp guard {ports port_list} {disabled-ports}
```

Description

Shows Simple Loop Protection Protocol (SLPP) Guard status for selected ports or all SLPP Guard-disabled ports.

Syntax Description

slpp	Specifies configuring SLPP.
guard	Specifies disabling a port as soon as an SLPP PDU is received.
ports	Specifies selecting ports that you want to see SLPP Guard status on.
<i>port_list</i>	Selects the ports that you want to see SLPP Guard status on.
disabled-ports	Specifies showing information about ports disabled by SLPP Guard.

Default

N/A.

Usage Guidelines

SLPP is an application that detects loops in a Split Multi-link Trunking (SMLT) network. SLPP Guard is a complementary feature that helps prevent loops in networks by administratively disabling an edge port if a switch receives an SLPP PDU from an SMLT network.

This command shows SLPP Guard status about selected ports (**ports** option) or shows all ports that have been disabled by SLPP Guard (**disabled-ports** option).

Example

The following example shows SLPP guard status for ports 15, 16, 31, 61, and 62

```
# show slpp guard ports 15 16 31 61 62
Port      SLPP Guard  State      Timeout    Age    Last Time Disabled
-----
15        Enabled     Disabled   300        46     Thu May 17 12:30:27 2018
16        Enabled     Monitoring 300        NA     NA
31        Enabled     Disabled   300        218    Thu May 17 12:10:27 2018
61        Enabled     Disabled   300        205    Thu May 17 12:20:27 2018
62        Enabled     Disabled   None       NA     Thu May 17 12:35:27 2018
=====
State:
  Disabled:    Port disabled due to SLPP PDU received.
  Monitoring:  Listening for SLPP PDU.
  None:       Port is down.
SLPP Guard Ethertype : 0xefef
```

The following example shows SLPP Guard-disabled ports:

```
# show slpp guard disabled-ports
Port      SLPP Guard  State      Timeout    Age    Last Time Disabled
-----
15        Enabled     Disabled   300        46     Thu May 17 12:30:27 2018
31        Enabled     Disabled   300        218    Thu May 17 12:10:27 2018
61        Enabled     Disabled   300        205    Thu May 17 12:20:27 2018
```

```
62          Enabled      Disabled      None      NA  Thu May 17 12:35:27 2018
=====
```

The following example shows SLPP Guard enabled by CLI:

```
# show slpp guard ports 1

Port      SLPP Guard  State      Timeout    Age Last Time Disabled  Enable
-----
1         Enabled    None       60         NA NA                    C-

=====

State:
    Disabled:    Port disabled due to SLPP PDU received.
    Monitoring:  Listening for SLPP PDU.
    None:        Port is down.
Enable Flags: (C) By CLI, (R) By Radius VSA

SLPP Guard Ethertype: 0X8102
```

The following example shows SLPP Guard enabled by Radius VSA:

```
# show slpp guard ports 1

Port      SLPP Guard  State      Timeout    Age Last Time Disabled  Enable
-----
1         Enabled    None       60         NA NA                    -R

=====

State:
    Disabled:    Port disabled due to SLPP PDU received.
    Monitoring:  Listening for SLPP PDU.
    None:        Port is down.
Enable Flags: (C) By CLI, (R) By Radius VSA

SLPP Guard Ethertype: 0X8102
```

The following example shows SLPP Guard enabled by Radius VSA and CLI:

```
# show slpp guard ports 1

Port      SLPP Guard  State      Timeout    Age Last Time Disabled  Enable
-----
1         Enabled    None       60         NA NA                    CR

=====

State:
    Disabled:    Port disabled due to SLPP PDU received.
    Monitoring:  Listening for SLPP PDU.
    None:        Port is down.
Enable Flags: (C) By CLI, (R) By Radius VSA

SLPP Guard Ethertype: 0X8102
```

History

This command was first available in ExtremeXOS 30.2.

Support for the Radius Attribute was introduced in ExtremeXOS 31.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show snmp

```
show snmp [get | get-next] object_identifier
```

Description

Displays the contents of an *SNMP* MIB object.

Syntax Description

<i>object_identifier</i>	Specifies the object identifier for an SNMP MIB object.
--------------------------	---

Default

N/A.

Usage Guidelines

Use the get option to establish an index into the SNMP MIB. After the get option is executed, you can use the get next option to step through the MIB objects.

Example

The following gets the contents of SNMP object 1.3.6.1.2.1.1.5.0:

```
show snmp get 1.3.6.1.2.1.1.5.0
system.5.0 = BD-12804
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show snmp notification-log entry

```
show snmp notification-log [ default | name hex hex_name ] entry
entry_index
```

Description

Displays a detailed summary of notification log configuration and status, as well as global configuration and status.

Syntax Description

default	The default log.
<i>name</i>	Specifies the name of the log.
<i>entry</i>	Specifies the index of an entry to be displayed in detail.
hex	Provide value in hexadecimal.
<i>hex_name</i>	Name of the log in hexadecimal.
<i>entry_index</i>	Specifies the entry index.

Default

Disabled.

Usage Guidelines

Use this command to display in detail the contents of a single entry from a notification log.

Example

The following example displays a detailed summary of *nmslog2* **entry** 452.

```
show snmp notification-log nmslog2 entry 452
Index      : 452
Date       : 05/09/2013
Time       : 08:45:12.41
Context    : VR-Default
OID        : 1.3.6.1.6.3.1.1.5.3

Variables:OID      : 1.3.6.1.2.1.1.3.0
Type       : Time Ticks
Value      : 86090700

OID        : 1.3.6.1.2.1.1.1.0
Type       : Octet String
Value      : 45:78:74:72:65:6d:65:58:4f
```

History

This command was first available in ExtremeXOS 15.5.

The **default** and **hex** keywords and *hex_name* variable were added in ExtremeXOS 15.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show snmp notification-log name

```
show snmp notification-log [default | name | hex hex_name]
```

Description

Displays a detailed summary of notification log configuration and status, as well as global configuration and status.

Syntax Description

default	The default log.
<i>name</i>	Specifies the name of the log.
hex	Provide value in hexadecimal.
<i>hex_name</i>	Name of the log in hexadecimal.

Default

Disabled.

Usage Guidelines

Use this command to display in detail the configuration and status of a notification log, as well as a summary of its contents.

Example

The following example displays a detailed summary of nmslog2.

```
show snmp notification-log nmslog2
Name                : nmslog2
Filter Profile      : filter1
Entry Limit         : 1500
Admin Status        : Disabled
Oper Status         : Admin Disabled
Entry Status        : Active
Storage Type        : Non Volatile
User                : nmsuser2
Security Model       : USM
Security Level       : Authentication Privacy
Notifications Logged: 1500
Notifications Bumped: 300
Entries             :
-----
      Index  Date      Time          Context          Notification OID
-----
          451 05/09/2013 07:00:01.17          1.3.6.1.6.3.1.1.5.1
          452 05/09/2013 08:45:12.41 VR-Default 1.3.6.1.6.3.1.1.5.3
          453 05/09/2013 09:10:05.32 VR-Default 1.3.6.1.6.3.1.1.5.4
-----

# show snmp notification-log default
Name                : default
```

```

Filter Profile      : test
Entry Limit        : System Managed
Admin Status       : Enabled
Oper Status        : Operational
Entry Status       : Active
Storage Type       : Non Volatile
User               : NA
Security Model     : NA
Security Level     : NA
Notifications Logged: 1
Notifications Bumped: 0
Entries           :

```

	Index	Date	Time	Context	Notification	OID
	1	06/11/2013	06:59:23.00		1.3.6.1.4.1.1916.0.15	
	2	06/11/2013	06:59:23.00		1.3.6.1.6.3.1.1.5.4	
	3	06/11/2013	06:59:25.00		1.3.6.1.6.3.1.1.5.4	

History

This command was first available in ExtremeXOS 15.5.

The **default** and **hex** keywords and *hex_name* variable were added in ExtremeXOS 15.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show snmp notification-log

```
show snmp notification-log
```

Description

Displays a summary of notification log configuration and status, as well as global configuration and status.

Syntax Description

This command has no keywords or variables.

Default

Disabled.

Usage Guidelines

Use this command to display a summary of notification log configuration and status, as well as global configuration and status.

Example

The following example displays log configuration and status:

```
show snmp notification-log
Global Entry Limit : 10000
System Managed Size : 1500
Global Age Out : None
Notifications Logged: 2100
Notifications Bumped: 300
Notification Logs :
-----
Name                               Filter Profile    Flags
Notifications
-----
                                     Logged    Bumped
-----
default                           all              EUAP      1500     300
nmslog1                           filter1          EFAV      0         0
nmslog2                           filter1          DMAN      600         0
-----
Flags: Admin Status: (D) Disable, (E) Enabled.
      Oper Status:  (F) No Filter, (M) Admin Disabled, (U) Up.
      Entry Status: (A) Active,   (I) Inactive, (T) Not Ready.
      Storage Type: (N) Non Volatile, (O) Other, (P) Permanent, (R) Read
Only,
                  (V) Volatile.
```

History

This command was first available in ExtremeXOS 15.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show snmp traps bfd

```
show snmp traps bfd
```

Description

This command displays session up/down trap reception for BFD.

Syntax Description

snmp	Configure <i>SNMP</i> specific settings.
traps	Configure SNMP Trap generation settings.
bfd	BFD-specific traps.

Default

Not applicable.

Usage Guidelines

Use this command to display SNMP trap reception for BFD session up/down.

Example

The following command displays SNMP Trap configuration for BFD:

```
#show snmp traps bfd
SNMP Traps for Session Down      : Enabled
SNMP Traps for Session Up       : Enabled
SNMP Traps Batch Delay           : 1000 ms
```

History

This command was first available in ExtremeXOS 15.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show snmp traps configuration

```
show snmp traps configuration
```

Description

Shows the current state of whether *SNMP* traps are set to occur when switch configurations are changed/saved.

Default

N/A

Example

The following example shows whether SNMP traps are set to occur when switch configurations are changed/saved:

```
# show snmp traps configuration
SNMP Traps for configuration change      : Disabled
SNMP Traps for configuration save       : Enabled
```

History

This command was first available in ExtremeXOS 22.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches..

show snmp vr_name

```
show snmp {vr} vr_name
```

Description

Displays the *SNMP* configuration and statistics on a virtual router.

Syntax Description

<i>vr_name</i>	Specifies the virtual router.
----------------	-------------------------------

Default

N/A.

Usage Guidelines

Use this command to display the SNMP configuration and statistics on a virtual router.

Example

The following command displays configuration and statistics for the virtual router *VR-Default*:

```
show snmp vr VR-Default
```

Following is sample output for the command:

```
x670.54 # show snmp vr "VR-Default"
SNMP access                : Enabled
SNMP ifMib ifAlias size    : Default
SNMP Traps                 : Enabled
SNMP TrapReceivers        :
  Destination              Source IP Address      Flags
  10.120.91.48 /162        32.1.0.2          2ET
  10.120.91.48 /162        32.2.0.2          2ET
  2001:2001:2001:2001:2001:2001:2001:2001 /16200 2001:2001:2001:2001:2001:2001:20022ET
  2001:2001:2001:2001:2001:2001:2004 /16200      2ET
  2334:0:0:0:0:0:2 /11555                          2ET
Flags:  Version: 1=v1 2=v2c 3=v3
        Mode: S=Standard E=Enhanced
        Notification Type: T=Trap I=Inform

SNMP stats:      InPkts 0      OutPkts 0      Errors 0      AuthErrors 0
                  Gets   0      GetNexts 0      Sets   0      Drops    0
SNMP traps:      Sent   0      AuthTraps Enabled
```

History

This command was first available in ExtremeXOS 12.4.2.

The SNMP ifMib ifAlias size status was added in ExtremeXOS 15.3.

The output was modified to display SNMPv3 notification targets and the notification type in ExtremeXOS 15.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show snmpv3 access

```
show snmpv3 access {[hex hex_group_name] | group_name]}
```

Description

Displays SNMPv3 access rights.

Syntax Description

hex	Specifies that the value to follow is to be supplied as a colon separated string of hex octets.
<i>group_name</i>	Specifies the name of the group to display.

Default

N/A.

Usage Guidelines

The `show snmpv3 access` command displays the access rights of a group. If you do not specify a group name, the command will display details for all the groups.

This command displays the SNMPv3 vacmAccessTable entries.

Example

The following command displays all the access details:

```
show snmpv3 access
```

The following is sample output from this command:

```
switch # show snmpv3 access
Group Name       : admin
Context Prefix   :
Security Model    : USM
Security Level    : Authentication Privacy
Context Match     : Exact
Read View        : defaultAdminView
Write View        : defaultAdminView
Notify View       : defaultNotifyView
Storage Type      : NonVolatile
```

```

Row Status      : Active
Group Name      : initial
Context Prefix  :
Security Model   : USM
Security Level   : No-Authentication No-Privacy
Context Match    : Exact
Read View       : defaultUserView
Write View      :
Notify View     : defaultNotifyView
Storage Type    : NonVolatile
Row Status      : Active
Group Name      : initial
Context Prefix  :
Security Model   : USM
Security Level   : Authentication No-Privacy
Context Match    : Exact
Read View       : defaultUserView
Write View      : defaultUserView
Notify View     : defaultNotifyView
Storage Type    : NonVolatile
Row Status      : Active
Group Name      : vlv2c_ro
Context Prefix  :
Security Model   : snmpv1
Security Level   : No-Authentication No-Privacy
Context Match    : Exact
Read View       : defaultUserView
Write View      :
Notify View     : defaultNotifyView
Storage Type    : NonVolatile
Row Status      : Active
Group Name      : vlv2c_ro
Context Prefix  :
Security Model   : snmpv2c
Security Level   : No-Authentication No-Privacy
Context Match    : Exact
Read View       : defaultUserView
Write View      :
Notify View     : defaultNotifyView
Storage Type    : NonVolatile
Row Status      : Active
Group Name      : vlv2c_rw
Context Prefix  :
Security Model   : snmpv1
Security Level   : No-Authentication No-Privacy
Context Match    : Exact
Read View       : defaultUserView
Write View      : defaultUserView
Notify View     : defaultNotifyView
Storage Type    : NonVolatile
Row Status      : Active
Group Name      : vlv2c_rw
Context Prefix  :
Security Model   : snmpv2c
Security Level   : No-Authentication No-Privacy
Context Match    : Exact
Read View       : defaultUserView
Write View      : defaultUserView
Notify View     : defaultNotifyView
Storage Type    : NonVolatile
Row Status      : Active
Group Name      : vlv2cNotifyGroup
Context Prefix  :
Security Model   : snmpv1

```

```
Security Level : No-Authentication No-Privacy
Context Match : Exact
Read View     :
Write View    :
Notify View   : defaultNotifyView
Storage Type  : NonVolatile
Row Status    : Active
Group Name    : vlv2cNotifyGroup
Context Prefix :
Security Model : snmpv2c
Security Level : No-Authentication No-Privacy
Context Match : Exact
Read View     :
Write View    :
Notify View   : defaultNotifyView
Storage Type  : NonVolatile
Row Status    : Active
Total num. of entries in vacmAccessTable : 9
```

The following command displays the access rights for the group group1:

```
show snmpv3 access group1
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show snmpv3 community

```
show snmpv3 community
```

Description

Displays information about *SNMP* community strings.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command displays information about and status of the SNMP community on the switch. This information is available to Administrator Accounts.

Example

The following command displays the community:

```
show snmpv3 community
```

The following is sample output from this command:

```
switch # show snmpv3 community
Community Index   : private
Community Name    : private
Security Name     : vlv2c_rw
Context EngineID  : 80:00:07:7c:03:00:04:96:27:b6:7b
Context Name      :
Transport Tag     :
Storage Type      : NonVolatile
Row Status        : Active
Community Index   : public
Community Name    : public
Security Name     : vlv2c_ro
Context EngineID  : 80:00:07:7c:03:00:04:96:27:b6:7b
Context Name      :
Transport Tag     :
Storage Type      : NonVolatile
Row Status        : Active
Total num. of entries in snmpCommunityTable : 2
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show snmpv3 context

```
show snmpv3 context
```

Description

Displays information about the SNMPv3 contexts on the switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command displays the entries in the View-based Access Control Model (VACM) context table (VACMContextTable).

Example

The following command displays information about the SNMPv3 contexts on the switch:

```
show snmpv3 context
```

The following is sample output from this command:

```
VACM Context Name :  
Note : This Version Supports one global context ("")
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show snmpv3 counters

```
show snmpv3 counters
```

Description

Displays SNMPv3 counters.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

The `show snmpv3 counters` command displays the following SNMPv3 counters:

- snmpUnknownSecurityModels.
- snmpInvalidMessages.
- snmpUnknownPDUHandlers.
- usmStatsUnsupportedSecLevels.

- usmStatsNotInTimeWindows.
- usmStatsUnknownUserNames.
- usmStatsUnknownEngineIDs.
- usmStatsWrongDigests.
- usmStatsDecryptionErrors.

Issuing the command `clear counters` resets all counters to zero.

Example

The following command displays all the SNMPv3 counters:

```
show snmpv3 counters
```

The following is sample output from this command:

```
snmpUnknownSecurityModels      : 0
snmpInvalidMessages            : 0
snmpUnknownPDUHandlers        : 0
usmStatsUnsupportedSecLevels    : 0
usmStatsNotInTimeWindows       : 0
usmStatsUnknownUserNames       : 0
usmStatsUnknownEngineIDs       : 0
usmStatsWrongDigests           : 0
usmStatsDecryptionErrors       : 0
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show snmpv3 engine-info

```
show snmpv3 engine-info
```

Description

Displays information about the SNMPv3 engine on the switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

The following show engine-info output is displayed:

- Engine-ID—Either the ID auto generated from MAC address of switch, or the ID manually configured.
- Engine Boots—Number of times the agent has been rebooted.
- Engine Time—Time since agent last rebooted, in centiseconds.
- Max. Message Size—Maximum *SNMP* Message size supported by the Engine (8192).

Example

The following command displays information about the SNMPv3 engine on the switch:

```
show snmpv3 engine-info
```

The following is sample output from this command:

```
SNMP Engine-ID      : 80:0:7:7c:3:0:30:48:41:ed:97 'H'
SNMP Engine Boots   : 1
SNMP Engine Time    : 866896
SNMP Max. Message Size : 8192
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show snmpv3 extreme-target-addr-ext

```
show snmpv3 extreme-target-addr-ext [[hex hex_addr_name] | addr_name]
```

Description

Displays information about SNMPv3 target addresses enhanced or standard mode.

Syntax Description

<i>hex_addr_name</i>	Specifies an identifier for the target address. The value is to be supplied as a colon separated string of hex octets.
<i>addr_name</i>	Specifies a string identifier for the target address.

Default

N/A.

Usage Guidelines

Use this command to display entries in the SNMPv3 extremeTargetAddressExtTable.

Example

The following command displays the entry for the target address named A1:

```
show snmpv3 extreme-target-addr-ext A1
```

The following is sample output from this command:

```
x670.26 # show snmpv3 extreme-target-addr-ext

Target Addr Name      : v1v2cNotifyTAddr1
Mode                  : Enhanced
IgnoreMPModel         : No
UseEventComm          : Yes

Target Addr Name      : v1v2cNotifyTAddr2
Mode                  : Enhanced
IgnoreMPModel         : No
UseEventComm          : Yes

Total num. of entries in extremeTargetAddrExtTable : 2
```

History

This command was first available in ExtremeXOS 10.1.

The hex_addr_name parameter was added in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show snmpv3 filter

```
show snmpv3 filter {[hex hex_profile_name] | profile_name} {{subtree}
object_identifier}
```

Description

Displays the filters that belong a filter profile.

Syntax Description

<i>hex_profile_name</i>	Specifies the filter profile to display. The value is to be supplied as a colon separated string of hex octets.
<i>profile_name</i>	Specifies the filter profile to display in ASCII format.
<i>object_identifier</i>	Specifies a MIB subtree.

Default

N/A.

Usage Guidelines

Use this command to display entries from the snmpNotifyFilterTable. If you specify a profile name and subtree, you will display only the entries with that profile name and subtree. If you specify only the profile name, you will display all entries for that profile name. If you do not specify a profile name, then all the entries are displayed.

Example

The following command displays the part of filter profile prof1 that includes the MIB subtree 1.3.6.1.4.1:

```
show snmpv3 filter prof1 subtree 1.3.6.1.4.1
```

The following is sample output from this command:

```
Profile Name      : prof1
Subtree           : 1.3.6.1.4.1
Mask              :
Type              : Included
Storage Type      : NonVolatile
Row Status        : Active
```

History

This command was first available in ExtremeXOS 10.1.

The hex_profile_name parameter was added in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show snmpv3 filter-profile

```
show snmpv3 filter-profile {[[hex hex_profile_name] | profile_name]}
                             {param [[hex hex_param_name] | param_name]}
```

Description

Displays the association between parameter names and filter profiles.

Syntax Description

<i>hex_profile_name</i>	Specifies the filter profile name. The value is to be supplied as a colon separated string of hex octets.
<i>profile_name</i>	Specifies the filter profile name in ASCII format.
<i>hex_param_name</i>	Specifies the parameter name. The values is to be supplied as a colon separated string of hex octets.
<i>param_name</i>	Specifies the parameter name in ASCII format.

Default

N/A.

Usage Guidelines

Use this command to display the snmpNotifyFilterProfileTable. This table associates a filter profile with a parameter name. The parameter name is associated with target addresses, and the filter profile is associated with a series of filters, so, in effect, you are associating a series of filters with a target address.

Example

The following command displays the entry with filter profile prof1 with the parameter name P1:

```
show snmpv3 filter-profile prof1 param P1
```

The following is sample output of this command:

```
Filter Profile Params Name : p1
Name                       : prof1
Storage Type               : NonVolatile
Row Status                 : Active
```

History

This command was first available in ExtremeXOS 10.1.

The hex_profile_name and hex_param_name parameters were added in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show snmpv3 group

```
show snmpv3 group {[hex hex_group_name] | group_name} {user [[hex
    hex_user_name] | user_name]}
```

Description

Displays the user name (security name) and security model association with a group name.

Syntax Description

<i>hex_group_name</i>	Specifies the group name to display. The value is to be supplied as a colon separated string of hex octets.
<i>group_name</i>	Specifies the group name to display. The value is to be supplied in ASCII format.
<i>hex_user_name</i>	Specifies the user name to display. The value is to be supplied as a colon separated string of hex octets.
<i>user_name</i>	Specifies the user name to display. The value is to be supplied in ASCII format.

Default

N/A.

Usage Guidelines

The `show snmpv3 group` command displays the details of a group with the given group name. If you do not specify a group name, the command will display details for all the groups.

This command displays the SNMPv3 vacmSecurityToGroupTable.

Example

The following command displays information about all groups for every security model and user name:

```
show snmpv3 group
```

The following is sample output from this command:

```
switch # sh snmpv3 group
Group Name       : v1v2c_ro
Security Name    : v1v2c_ro
Security Model   : snmpv1
Storage Type     : NonVolatile
Row Status      : Active
Group Name       : v1v2c_rw
Security Name    : v1v2c_rw
Security Model   : snmpv1
Storage Type     : NonVolatile
Row Status      : Active
Group Name       : v1v2c_ro
Security Name    : v1v2c_ro
Security Model   : snmpv2c
Storage Type     : NonVolatile
Row Status      : Active
Group Name       : v1v2c_rw
Security Name    : v1v2c_rw
Security Model   : snmpv2c
```



```

Storage Type      : NonVolatile
Row Status       : Active
Group Name       : admin
Security Name    : admin
Security Model   : USM
Storage Type     : NonVolatile
Row Status      : Active
Group Name      : initial
Security Name   : initial
Security Model  : USM
Storage Type    : NonVolatile
Row Status     : Active
Group Name     : initial
Security Name  : initialmd5
Security Model : USM
Storage Type   : NonVolatile
Row Status    : Active
Group Name    : initial
Security Name : initialsha
Security Model : USM
Storage Type  : NonVolatile
Row Status   : Active
Group Name   : initial
Security Name : initialmd5Priv
Security Model : USM
Storage Type : NonVolatile
Row Status  : Active
Group Name  : initial
Security Name : initialshaPriv
Security Model : USM
Storage Type : NonVolatile
Row Status  : Active
Total num. of entries in vacmSecurityToGroupTable : 10

```

The following command shows information about the group testgroup and user name testuser:

```
show snmpv3 group testgroup user testuser
```

The following is sample output from this command:

```

Group Name      : testgroup
Security Name   : testuser
Security Model  : USM
Storage Type    : NonVolatile
Row Status     : Active

```

History

This command was first available in ExtremeXOS 10.1.

The hex_group_name and hex_user_name parameters were added in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show snmpv3 mib-view

```
show snmpv3 mib-view {[[hex hex_view_name] | view_name] {subtree
object_identifier}}
```

Description

Displays a MIB view.

Syntax Description

<i>hex_view_name</i>	Specifies the name of the MIB view to display. The value is to be supplied as a colon separated string of hex octets.
<i>view_name</i>	Specifies the name of the MIB view to display. The value is to be supplied in ASCII format.
<i>object_identifier</i>	Specifies the object identifier of the view to display.

Default

N/A.

Usage Guidelines

The [show snmpv3 mib-view](#) command displays a MIB view. If you do not specify a view name, the command will display details for all the MIB views. If a subtree is not specified, then all subtrees belonging to the view name will be displayed.

This command displays the SNMPv3 vacmViewTreeFamilyTable.

Example

The following command displays all the view details:

```
show snmpv3 mib-view
```

The following is sample output from this command:

```
switch # sh snmpv3 mib-view
View Name       : defaultUserView
MIB Subtree     : 1
Mask            :
View Type       : Included
Storage Type    : NonVolatile
Row Status      : Active
View Name       : defaultUserView
MIB Subtree     : 1.3.6.1.6.3.16
Mask            :
View Type       : Excluded
Storage Type    : NonVolatile
Row Status      : Active
View Name       : defaultUserView
```

```

MIB Subtree      : 1.3.6.1.6.3.18
Mask             :
View Type        : Excluded
Storage Type     : NonVolatile
Row Status       : Active
View Name        : defaultUserView
MIB Subtree      : 1.3.6.1.6.3.15.1.2.2.1.4
Mask             :
View Type        : Excluded
Storage Type     : NonVolatile
Row Status       : Active
View Name        : defaultUserView
MIB Subtree      : 1.3.6.1.6.3.15.1.2.2.1.6
Mask             :
View Type        : Excluded
Storage Type     : NonVolatile
Row Status       : Active
View Name        : defaultUserView
MIB Subtree      : 1.3.6.1.6.3.15.1.2.2.1.9
Mask             :
View Type        : Excluded
Storage Type     : NonVolatile
Row Status       : Active
View Name        : defaultAdminView
MIB Subtree      : 1
Mask             :
View Type        : Included
Storage Type     : NonVolatile
Row Status       : Active
View Name        : defaultNotifyView
MIB Subtree      : 1
Mask             :
View Type        : Included
Storage Type     : NonVolatile
Row Status       : Active
Total num. of entries in vacmViewTreeFamilyTable : 8

```

The following command displays a view with the view name Rview and subtree 1.3.6.1.2.1.1:

```
show snmpv3 mib-view Rview subtree 1.3.6.1.2.1.1
```

History

This command was first available in ExtremeXOS 10.1.

The `hex_view_name` parameter was added in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show snmpv3 notify

This command displays the `snmpNotifyTable`.

```
show snmpv3 notify {[[hex hex_notify_name] | notify_name]}
```

Description

Displays the notifications that are set.

Syntax Description

<i>hex_notify_name</i>	Specifies the parameter name associated with the target. The value is to be supplied as a colon separated string of hex octets.
<i>notify_name</i>	Specifies the parameter name associated with the target. The value is to be supplied in ASCII format.

Default

N/A.

Usage Guidelines

Use this command to display entries from the SNMPv3 snmpNotifyTable. This table lists the notify tags that the agent will use to send notifications (traps).

If no notify name is specified, all the entries are displayed.

Example

The following command displays the notify table entries:

```
show snmpv3 notify
```

The following is sample output from this command:

```
# show snmpv3 notify
Notify Name      : defaultNotify
Tag              : defaultNotify
Type             : Trap
Storage Type     : NonVolatile
Row Status       : Active
Notify Name      : xyz
Tag              : xyz1
Type             : Inform
Storage Type     : NonVolatile
Row Status       : Active
Total entries in snmpNotifyTable : 2
```

History

This command was first available in ExtremeXOS 10.1.

The hex_notify_name parameter was added in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show snmpv3 target-addr

```
show snmpv3 target-addr {[hex hex_addr_name] | addr_name}
```

Description

Displays information about SNMPv3 target addresses.

Syntax Description

<i>hex_addr_name</i>	Specifies an identifier for the target address. The value is to be supplied as a colon separated string of hex octets.
<i>addr_name</i>	Specifies a string identifier for the target address.

Default

N/A.

Usage Guidelines

Use this command to display entries in the SNMPv3 snmpTargetAddressTable. If no target address is specified, the entries for all the target addresses will be displayed.

To view the source IP address, use the [show management](#) command.

Example

The following command displays the entry for the target address named A1:

```
show snmpv3 target-addr A1
```

The following is sample output from this command:

```
x670.25 # show snmpv3 target-addr

Target Addr Name      : v1v2cNotifyTAddr1
TDomain               : 1.3.6.1.6.1.1
TAddress              : 10.120.91.48, 162
TMask                 :
Timeout               : 1500
Retry Count           : 3
Tag List              : defaultNotify
Params               : v1v2cNotifyParam1
Storage Type          : NonVolatile
Row Status             : Active

Target Addr Name      : v1v2cNotifyTAddr2
TDomain               : 1.3.6.1.6.1.1
TAddress              : 2001:0:0:0:0:0:5, 162
```

```

TMask          :
Timeout        : 1500
Retry Count    : 3
Tag List       : defaultNotify
Params         : v1v2cNotifyParam2
Storage Type   : NonVolatile
Row Status     : Active

Total num. of entries in snmpTargetAddrTable : 2

```

History

This command was first available in ExtremeXOS 10.1.

This command was modified in ExtremeXOS 11.0 to display a list of tags if more than one was configured and to display the timeout value for the entry in the snmpTargetAddrTable. This command was also modified to support the hex_addr_name parameter.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show snmpv3 target-params

```
show snmpv3 target-params {[hex hex_target_params] | target_params]}
```

Description

Displays the information about the options associated with the parameter name.

Syntax Description

<i>hex_target_params</i>	Specifies the parameter to display. The value is to be supplied as a colon separated string of hex octets.
<i>target_params</i>	Specifies the parameter name to display. The value is to be supplied in ASCII format.

Default

N/A.

Usage Guidelines

Use this command to display entries from the SNMPv3 snmpTargetParamsTable. This table specifies the message processing model, security level, security model, and the storage parameters for messages to any target addresses associated with a particular parameter name.

If no parameter name is specified, all the entries are displayed.

Example

The following command displays the target parameter entry named P1:

```
show snmpv3 target-params P1
```

The following is sample output from this command:

```
Target Params Name      : p1
MP Model                : snmpv2c
Security Model          : snmpv2c
User Name               : testuser
Security Level          : No-Authentication No-Privacy
Storage Type            : NonVolatile
Row Status              : Active
```

History

This command was first available in ExtremeXOS 10.1.

The hex_target_params parameter was added in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show snmpv3 user

```
show snmpv3 user { [[hex hex_user_name] | user_name] }
```

Description

Displays detailed information about the user.

Syntax Description

<i>hex_user_name</i>	Specifies the user name to display. The value is to be supplied as a colon separated string of hex octets.
<i>user_name</i>	Specifies the user name to display. The value is to be supplied in ASCII format.

Default

N/A.

Usage Guidelines

The `show snmpv3 user` command displays the details of a user. If you do not specify a user name, the command will display details for all the users. The authentication and privacy passwords and keys will not be displayed.

The user entries in SNMPv3 are stored in the USMUserTable, so the entries are indexed by EngineID and user name.

Example

The following command lists all user entries:

```
show snmpv3 user
```

The following is sample output from this command:

```
switch # sh snmpv3 user
Engine-ID      : 80:00:07:7c:03:00:04:96:27:b6:7b 'H'
User Name      : admin
Security Name   : admin
Authentication  : HMAC-MD5
Privacy        : DES
Storage Type    : NonVolatile
Row Status     : Active
Engine-ID      : 80:00:07:7c:03:00:04:96:27:b6:7b 'H'
User Name      : initial
Security Name   : initial
Authentication  : No-Authentication
Privacy        : No-Privacy
Storage Type    : NonVolatile
Row Status     : Active
Engine-ID      : 80:00:07:7c:03:00:04:96:27:b6:7b 'H'
User Name      : initialmd5
Security Name   : initialmd5
Authentication  : HMAC-MD5
Privacy        : No-Privacy
Storage Type    : NonVolatile
Row Status     : Active
Engine-ID      : 80:00:07:7c:03:00:04:96:27:b6:7b 'H'
User Name      : initialsha
Security Name   : initialsha
Authentication  : HMAC-SHA
Privacy        : No-Privacy
Storage Type    : NonVolatile
Row Status     : Active
Engine-ID      : 80:00:07:7c:03:00:04:96:27:b6:7b 'H'
User Name      : initialmd5Priv
Security Name   : initialmd5Priv
Authentication  : HMAC-MD5
Privacy        : DES
Storage Type    : NonVolatile
Row Status     : Active
Engine-ID      : 80:00:07:7c:03:00:04:96:27:b6:7b 'H'
User Name      : initialshaPriv
Security Name   : initialshaPriv
Authentication  : HMAC-SHA
Privacy        : DES
Storage Type    : NonVolatile
Row Status     : Active
Total num. of entries in usmUserTable : 6
```

The following command lists details for the specified user, testuser:

```
show snmpv3 user testuser
```


History

This command was first available in ExtremeXOS 10.1.

The `hex_user_name` parameter was added in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

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 SNTP configuration:

```
show sntp-client
```

The following is sample output from this command:

```
SNTP client is enabled
SNTP time is valid
Primary server: 172.17.1.104
Secondary server: 172.17.1.104
Query interval: 64
Last valid SNTP update: From server 172.17.1.104, on Wed Oct 30 22:46:03 2003
SNTPC Statistics:
Packets transmitted:
to primary server:      1
to secondary server:    0
Packets received with valid time:
from Primary server:    1
from Secondary server:  0
from Broadcast server:  0
```

```
Packets received without valid time:
from Primary server:      0
from Secondary server:    0
from Broadcast server:    0
Replies not received to requests:
from Primary server:      0
from Secondary server:    0
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show ssh2

```
show ssh2
```

Description

Shows all SSHv2 configuration information, including the ciphers/MACs that are enabled, and Diffie-Hellman minimal supported group.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Example

The following example shows all SSHv2 configuration information:

```
SSH module configuration details:
SSH Access      : Disabled
Key validity    : Invalid
Key type        : RSA 2048
TCP port        : 22
VR              : all
Access profile   : not set
Secure Mode     : Off
Diffie-Hellman Groups : 18 (8192 bits)
Max Auth Tries  : 3
Idle time       : 60 minutes
Rekey Interval   : 4096 MB and no time limit
Ciphers         : aes128-cbc, 3des-cbc, aes192-cbc, aes256-cbc, rijndael-
cbc@lysator.liu.se, aes128-ctr, aes192-ctr, aes256-ctr, chacha20-poly1305@openssh.com
Macs            : hmac-md5-etm@openssh.com, hmac-sha1-etm@openssh.com, hmac-
sha2-256-etm@openssh.com, hmac-sha2-512-etm@openssh.com, hmac-sha1-96-etm@openssh.com,
hmac-md5-96-etm@openssh.com, hmac-md5, hmac-sha1, hmac-sha2-256, hmac-sha2-512, hmac-
```

```

sha1-96, hmac-md5-96
Public key algorithms : ssh-rsa, ssh-dss, x509v3-sign-rsa, x509v3-sign-dss
Login grace timeout   : 100 seconds

```

The following command displays x509v3 OSCP attributes (lines 18-22):

```

# show ssh2
SSH module configuration details:
SSH Access           : Disabled
Key validity         : Invalid
Key type             : none
TCP port            : 22
VR                  : all
Access profile       : not set
Secure Mode          : Off
Diffie-Hellman Groups : 14 (2048 bits), 16 (4096 bits), 18 (8192 bits)
Max Auth Tries       : 3
Idle time            : 60 minutes
Rekey Interval       : 4096 MB and no time limit
Ciphers              : chacha20-poly1305@openssh.com, aes128-ctr, aes192-ctr, aes256-ctr
Macs                 : hmac-sha2-256-etm@openssh.com, hmac-sha2-512-etm@openssh.com,
hmac-sha1-etm@openssh.com, hmac-sha2-256, hmac-sha2-512, hmac-sha1
Public key algorithms : ssh-rsa, x509v3-sign-rsa, x509v3-sign-dss
Login grace timeout   : 120 seconds
X509v3 OSCP Attributes:
  OSCP                : On
  Nonce               : On
  Signer ocsf-nocheck : On
  Override Server URL : http://sshocsp:2023

```

The following command displays x509v3 RADIUS Authentication (lines 18-22):

```

# show ssh2
SSH module configuration details:
SSH Access           : Enabled
Key validity         : Valid
Key type             : RSA 2048
TCP port            : 22
VR                  : all
Access profile       : not set
Secure Mode          : Off
Diffie-Hellman Groups : 14 (2048 bits), 16 (4096 bits), 18 (8192 bits)
Max Auth Tries       : 3
Idle time            : 60 minutes
Rekey Interval       : 4096 MB and no time limit
Ciphers              : chacha20-poly1305@openssh.com, aes128-ctr, aes192-ctr,
aes256-ctr
Macs                 : hmac-sha2-256-etm@openssh.com, hmac-sha2-512-
etm@openssh.com, hmac-sha1-etm@openssh.com, hmac-sha2-256, hmac-sha2-512, hmac-sha1
Public key algorithms : ssh-rsa, x509v3-sign-rsa, x509v3-sign-dss
Login grace timeout   : 120 seconds
X509v3 RADIUS Authentication :
  Password authentication : On
  Username overwrite       : On
  Username strip domain    : On
  Username use domain      : abcdef.com

```

History

This command was first available in ExtremeXOS 22.1.

Information about rekey interval and public key algorithms was first available in ExtremeXOS 22.3.

Information about key type was added in ExtremeXOS 22.5.

Information about the login grace timeout period was added in ExtremeXOS 30.7.

x509v3 OCSP attributes and RADIUS Authentication were added in ExtremeXOS 32.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show ssh2 ciphers macs

```
show ssh2 {ciphers | macs}
```

Description

Displays configured SSHv2 ciphers and Message Authentication Codes (MACs).

Syntax Description

ciphers	Shows configured ciphers.
macs	Shows configured MACs.

Default

N/A.

Example

The following example shows all MACs that are enabled:

```
# show ssh2 macs
Macs: hmac-md5-etm@openssh.com, hmac-sha1-etm@openssh.com,
hmac-sha2-256-etm@openssh.com, hmac-sha2-512-etm@openssh.com, hmac-sha1-96-
etm@openssh.com,
hmac-md5-96-etm@openssh.com, hmac-md5, hmac-sha1, hmac-sha2-256,
hmac-sha2-512, hmac-sha1-96, hmac-md5-96
```

History

This command was first available in ExtremeXOS 22.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show ssh2 private-key

```
show ssh2 private-key
```

Description

Displays the ssh2 server's private key.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command displays the ssh server's private key which can be used to configure the key later or on another switch by using the `configure ssh2 key {pregenerated}` command. The key is saved in the switch's EEPROM.

To erase the key from the EEPROM, use the `unconfigure switch` command.

History

This command was first available in ExtremeXOS 12.1.

This command was added to ExtremeXOS 11.6 SR, and ExtremeXOS 12.0 SR.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show sshd2 user-key

```
show sshd2 user-key {key_name {users}}
```

Description

Displays the user names bound to a key.

Syntax Description

<i>key_name</i>	Specifies the name of the public key.
users	Specifies the name of the users.

Default

N/A.

Usage Guidelines

This command displays the names of the users that are bound to a public key.

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show ssl

```
show ssl {[trusted-ca | ocsp-signature-ca] [file_name | all]}
        {manufacturing}{certificate | detail}
```

Description

Displays the secure socket layer (SSL) configuration.

Syntax Description

manufacturing	Shows manufacturer-installed certificates.
trusted-ca	Shows trusted CA certificates.
ocsp-signature-ca	Shows OCSP signature CA certificates.
<i>file_name</i>	Prints specified certificate file.
all	Prints all certificates.
certificate	Prints the certificate.
detail	Displays detailed SSL information.

Default

N/A.

Usage Guidelines

This command displays the following information:

- HTTPS port configured. This is the port on which the clients will connect.
- Length of the RSA key (the number of bits used to generate the private key).
- Basic information about the stored certificate.

Example

The `show ssl` command displays the SSL configuration. The following is sample output from this command:

```
HTTPS Port Number: 443
Private Key matches with the Public Key in certificate. (or Private key does not match
with the Public Key in the certificate)
RSA Key Length: 1024
Certificate:
Data:
Version: 1 (0x0)
Serial Number: 6 (0x6)
Signature Algorithm: md5WithRSAEncryption
Issuer: C=AU, O=CryptSoft Pty Ltd, CN=Test CA (1024 bit)
Validity
Not Before: Oct 16 22:31:03 2000 GMT
Not After : Jan 14 22:31:03 2003 GMT
Subject: C=AU, O=CryptSoft Pty Ltd, CN=Server test cert (512 bit)
Manufacturing certificate: Not present
```

History

This command was first available in the ExtremeXOS 11.2.

The **trusted-ca** and **ocsp-signature-ca** options were added in ExtremeXOS 22.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show ssl csr

```
show ssl csr
```

Description

Shows the generated certificate signing request (CSR).

Syntax Description

ssl	Specifies SSL (Secure Sockets Layer).
csr	Specifies showing the CSR (certificate signing request).

Default

N/A.

Example

The following example shows the generated CSR:

```
# show ssl csr
-----BEGIN CERTIFICATE REQUEST-----
MIIC3TCCAcUCAQIwgZcxCzAJBgNVBAYTA1VTMQ0wCwYDVQQKDARFWFRSMREwDwYD
VQQDDAhjc3JfdGVzdDEXMBUGA1UECAwOTm9ydGggQ2Fyb2xpbmExEDAOBgNVBACM
B1JhbGVpZ2gxDDAKBgNVBAsMA1JEVTEtMCsGCSqGSIb3DQEJARYebHBldHR5am9o
bkBleHRyZW1lbmV0d29ya3MuY29tMIIBIjANBgkqhkiG9w0BAQEFAAOCAQ8AMIIB
CgKCAQEAm43c60n1XXkk1MMvK+ovX8fAhWRu8j7TAKGrSENqEhmS0BI05bjZLsj/
loulgsPXQA17W4010OMt5w9zcMCNmSf47PJwpQZpo4msAW8uSp7IMM9Ctv0a8oLr
kArzh3F+Gp0cAe7LycOthiXINKKwmzWpNwHmGbrwAhbd3grShurvUU7n0b+1Xcle
YH5J/HnGq+j6Lb+iNF2RbCactChF0aeT7DKXZaIt8s+p9ib3XQXUNvGoP+4M/Eoq
dHfOwppvBJeL3EyhjkEmz456nwdtsY8deNi/ssW+VJJWpGPONNLo+1lwD7BksCptJ
Pf20atDCFj6bFAo6N9gbdkhldI3euwIDAQABoAAwDQYJKoZIhvcNAQENBQADggEB
AikoEBWhrPmL4tf0KSgKeadfODJ6Nipkcyof9YZ9AceJhtgMmBFmMfcUrE+3e28j
asXQpEc5hLkc8fyRMNjDHuuz2d6uWju+K/TqVNT094bvbvySFsdBKjLcOADlRP0m
CIMCCiAiaFhtmLE5Sg6BoYctJ2jRNJ4UQOejeclcg80+qaXu6u7xAg5emGMtJizE
bvePhgSdhYTCFGngFrg3pZXHHTvRB7t54oYGG7yYdFb3jyW8CzckxnkiTV87fxHP
ojUeAwXet1AfI8cof1Dfmf6gKnBLMzrz5DMDmqdJgE2HgLLZCLv+JZbjbmowLrDL
DhG3F97QQkwROTpJfmrSsaU=
-----END CERTIFICATE REQUEST-----
```

History

This command was first available in ExtremeXOS 22.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show stack-ports debounce

```
show stack-ports {port-list} debounce
```

Description

This command displays the current debounce time configured in stack-ports.

Syntax Description

<i>port_list</i>	Specifies one or more stacking ports.
------------------	---------------------------------------

Default

N/A.

Usage Guidelines

Use this command to view the current debounce time configured in stack-ports. Specifying the stack-port allows to view the debounce time for particular stack-port alone.

Example

The following example displays the output of the `show stack-ports 1:1 1:2 debounce` command:

Stack Port	Debounce Time (ms)
1:1	0
1:2	0

History

This command was first available in ExtremeXOS 15.3.4.

Platform Availability

This command is available on all stackable switches.

show stacking

```
show stacking
```

Description

The `show stacking` command shows a summary of the nodes in the stack topology.

The `show stacking` command shows all nodes that are in the stack topology.

Syntax Description

This command has no arguments or variables.

Default

There is no default value for this command.

Usage Guidelines

The asterisk (*) that precedes the node MAC address indicates the node on which this command is being executed, that is, the node to which the user is logged in.

The node MAC address is the address that is factory assigned to the stackable.

The slot number shown is the number currently in use by the related node. Since slot number configuration only takes effect during node initialization, a change in configured value alone does not cause a change to the slot number that is in use. Slot numbers show as hyphen (-) characters on nodes that have stacking disabled.

The Stack State shows the state values.

The Role is one of the following: Master, Backup, Standby, or *none*.

In a ring topology, the node on which this command is executed is always the first node displayed. The order of the nodes shown in the display is the order of their physical connection in the ring.

Even though the stack topology can be a ring, the active topology can simultaneously be a daisy chain because it is only a proper subset of the stack topology. If the node on which this command is executed is not active, the line

```
Active Topology is a ____
```

is replaced by the line

```
This node is not in an Active Topology.
```

The daisy chain topology is displayed in the order of physical connection. The master node detects the two nodes in the stack topology that have only one operating link, and these nodes become the ends of the stack. Such nodes always display at the top and bottom of the output.

It is possible for a node to be in Stabilizing or Waiting state and still be in the active topology. This is because it is possible for an active node to move to these states when a topology change is detected. Once a node becomes active, the node remains an active node until it reboots or an overflow condition occurs.

The Flags have the following definitions:

- The C flag indicates that the related node is a candidate for membership of the same active topology to which the node on which the command is executed would belong.
- The A flag indicates that the related node is an active node in the active topology of which the node on which the command is run is also a candidate node. Being an active node is necessary but not sufficient for presence of the node in a slot. Once the node has fully initialized, the active node appears as Present in the show slot display.
- The O flag indicates that the related node is probably an active node in an active topology for which the node on which this command is being run is not a candidate.

The O flag is useful for the case where there is an inhibited link or a disabled or failed node that separates two active topologies. One active topology may contain the local node, and all other nodes in this active topology do not have the O flag set. All nodes that are members of an active topology that is separated by an inhibited link from the active topology that contains the local node have only the O flag set. All possibly active nodes have the O flag set if the local node is not a member of any active topology. For any node for which the O flag is set, the C and A flags are not set and vice-versa.

The following information is displayed:

- Stack Topology is a ring or daisy-chain.
- Active Topology is a ring or daisy-chain (or This node is not in an Active Topology.).
- For each node:
 - Node MAC address (factory assigned).
 - Slot number in use.

- Stack State:
 - Disabled - Node is not configured for stacking.
 - Failed - Node can't come up in the stack because it has a duplicate slot number.
 - Overflow - The node has detected that there are more nodes in the stack topology than are allowed.
 - Listening - Initial state when attempting to join the stack. The node is checking to see if its configured slot number duplicates that of another node. The node cannot be an active node in this state.
 - Stabilizing - Node is waiting until it sees no new topology changes. The node may or may not be an active node in this state.
 - Waiting - Topology has stabilized, if the active topology is to be a ring, and stacking link blocking is being performed. The node may or may not be an active node in this state.
 - Active - The node is an active node and is fully programmed to operate in the active topology.
- Node role (master, backup, standby, or other transient node state).
- Flags describing the node's membership in the active topology.
- Whether or not the node is this node, that is, the node on which the command is run.

Example

The following example shows the output of show stacking command:

```
Slot-1 Stack.30 # show stacking
Stack Topology is a Ring
Active Topology is a Daisy-Chain
Node MAC Address      Slot  Stack State  Role      Flags
-----
*00:04:96:26:60:DD    1      Stabilizing  Master    CA-
00:04:96:26:60:EE    4      Stabilizing  Standby   C--
00:04:96:26:60:FF    -      Disabled    Master    ---
00:04:96:26:60:AA    -      Disabled    Master    ---
00:04:96:26:60:88    -      Disabled    Master    ---
00:04:96:26:60:99    -      Disabled    Master    ---
00:04:96:26:60:BB    2      Stabilizing  Standby   C--
00:04:96:26:60:CC    3      Active      Backup    CA-
(*) Indicates This Node
Flags: (C) Candidate for this active topology, (A) Active node,
(O) node may be in Other active topology
Slot-1 Stack.31 #
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available with all licenses and platforms that support the SummitStack feature. For information about which licenses and platforms support the SummitStack feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

show stacking configuration

show stacking configuration

Description

Shows how the nodes are configured in a stack topology. The configured values shown are the ones actually stored in the remote nodes at the time you issue this command.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Some stacking parameters do not take effect until the next restart, the configured values and the values currently being used are both shown. Specifically, this applies to the slot number, whether or not stacking is enabled, the master-capable configuration, the license level restriction, and the stack MAC configuration.

The only parameters that take effect without a reboot are the node priority and the alternate management IP subnetwork and gateway.

The Stack MAC in use line can display the following values:

- If the command is executed on the master node:
 - *none* if there is no stack MAC configured.
 - The stack MAC configured on the master node.
- If the command is executed on a non-master node:
 - *unknown*. The stack MAC address is only known by the executing master node. In this case, the M and m flags are not set. The i flag is set if there is a stack MAC configured locally.

Identified with the asterisk, the current node is the one on which the `show stacking configuration` command is executed.

A node identified with the ? character indicates that timely attempts to fetch the configuration information from the node have failed. There are two possible reasons for this display:

- Communications with the node have been lost, in which case the node will probably be removed from the stack topology shortly.
- The node is too busy to respond in time.

A row that displays the ? indicator shows the last values that were received from the node. If no values were ever received, all configured values show as not configured (-) or *none*. The node MAC address and the slot number that is currently in use are still displayed.

Example

The following example shows the stacking configuration for a switch:

```
Slot-1 Stack.2 # show stacking configuration
Stack MAC in use: 02:04:96:26:6b:ed
Node          Slot      Alternate
MAC Address   Cfg Cur Prio Mgmt IP / Mask  Alternate
-----
*00:04:96:26:6b:ed 1    1    Auto <none>      <none>      CcEeMm--- --
00:04:96:34:d0:b8 2    2    Auto <none>      <none>      CcEeMm--- --
* - Indicates this node
Flags: (C) master-Capable in use, (c) master-capable is configured,
(E) Stacking is currently Enabled, (e) Stacking is configured Enabled,
(M) Stack MAC in use, (m) Stack MACs configured and in use are the same,
(N) Stack link protocol Enhanced in use, (n) Stack link protocol Enhanced configured,
(i) Stack MACs configured and in use are not the same or unknown,
(-) Not in use or not configured
License level restrictions: (C) Core, (A) Advanced edge, or (E) Edge in use,
(c) Core, (a) Advanced edge, or (e) Edge configured,
(-) Not in use or not configured
```

The following example shows the stacking configuration for a ExtremeSwitching 5520 series switch:

```
# show stacking configuration
Stack MAC in use: 02:04:96:f1:b8:00
Node          Slot      Alternate
MAC Address   Cfg Cur Prio Mgmt IP / Mask  Alternate
-----
*00:04:96:f1:b8:00 1    -    Auto <none>      <none>      -c-----iNn -b
* - Indicates this node
Flags: (C) master-Capable in use, (c) master-capable is configured,
(E) Stacking is currently Enabled, (e) Stacking is configured Enabled,
(M) Stack MAC in use, (m) Stack MACs configured and in use are the same,
(i) Stack MACs configured and in use are not the same or unknown,
(N) Enhanced protocol is in use, (n) Enhanced protocol is configured,
(-) Not in use or not configured
License level restrictions: (B) Base, or (P) Premier in use,
(b) Base, or (p) Premier configured,
(-) Not in use or not configured
```

History

This command was first available in ExtremeXOS 12.0.

Licensing information specific to ExtremeSwitching 5520 series switches was added in ExtremeXOS 31.1.

Platform Availability

This command is available with all licenses and platforms that support the SummitStack feature. For information about which licenses and platforms support the SummitStack feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

show stacking detail

```
show stacking {node-address node_address | slot slot_number} detail
```

Description

This command displays information about a specified node.

Syntax Description

<i>node_address</i>	Specifies the MAC address of a node in the stack. To view the MAC addresses for all nodes in a stack, enter the show stacking command.
<i>slot_number</i>	Specifies the slot number of the target node. To view the slot numbers, enter the show stacking command.

Default

N/A.

Usage Guidelines

If no node is specified, the output is generated for all nodes in the stack topology. If the specified node does not exist, an error message appears. The slot parameter is available only for active nodes in the same active topology as the node on which the command is run. The node-address parameter is always available.

Current information represents stacking states and configured values that are currently in effect. Configured information is that which takes effect at node reboot only. Thus, differences between values in use and values configured can be seen here. The advantages of this command over the [show stacking configuration](#) command is that the values in use and the configured values are fully expanded without the need for flags. You can also see the port state information of the node(s).

The roles values are: Master, Backup, Standby, and *none*.

License level restrictions can be Edge, Advanced Edge, or Core.

If one of the fields in the example below is missing on your switch, your switch does not support the feature that the field represents.

Example

The following is a sample output of this command:

```
Slot-1 Stack.33 # show stacking slot 1 detail
Stacking Node 00:04:96:26:6b:ec information:
Current:
Stacking           : Enabled
Role               : Master
Priority           : Automatic
Slot number        : 1
Stack state        : Active
Master capable?    : Yes
Stacking protocol  : Enhanced
License level restriction : <none>
In active topology? : Yes
Factory MAC address : 00:04:96:26:6b:ec
```

```

Stack MAC address      : 02:04:96:26:6b:ec
Alternate IP address   : <none>
Alternate gateway      : <none>
Stack Port 1:
State                  : Operational
Blocked?              : No
Control path active?   : Yes
Selection              : Alternate (23)
Stack Port 2:
State                  : Operational
Blocked?              : Yes
Control path active?   : Yes
Selection              : Native
Configured:
Stacking               : Enabled
Master capable?        : Yes
Slot number            : 1
Stack MAC address      : 02:04:96:26:6b:ec
Stacking protocol      : Enhanced
License level restriction : <none>
Stack Port 1:
Selection              : Alternate (23)
Stack Port 2:
Selection              : Native

```

The following example shows the stacking information for a ExtremeSwitching 5520 series switch node:

```

# show stacking detail
Stacking Node 00:04:96:f1:b8:00 information:
  Current:
    Stacking           : Disabled
    Role               : Master
    Priority            : Automatic
    Slot number        : N/A
    Stack state         : Disabled
    Master capable?     : N/A
    Stacking protocol   : Enhanced
    License level restriction : Base
    In active topology? : No
    Factory MAC address : 00:04:96:f1:b8:00
    Stack MAC address   : N/A
    Alternate IP address : <none>
    Alternate gateway   : <none>
    Stack Port 1:
      State             : Link down
      Blocked?          : No
      Control path active? : No
    Stack Port 2:
      State             : Link Down
      Blocked?          : No
      Control path active? : No
  Configured:
    Stacking           : Disabled
    Master capable?     : Yes
    Slot number        : 1
    Stack MAC address   : 02:04:96:f1:b8:00
    Stacking protocol   : Enhanced
    License level restriction : Base

```

History

This command was first available in ExtremeSwitching 12.0.

The Stacking protocol and Stack Port Selection fields were added in ExtremeSwitching 12.5.

Licensing information specific to ExtremeSwitching 5520 series switches was added in ExtremeXOS 31.1.

Platform Availability

This command is available with all licenses and platforms that support the SummitStack feature. For information about which licenses and platforms support the SummitStack feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

show stacking stack-ports

```
show stacking stack-ports
```

Description

This command displays the port states of each node in the stack topology and the connections between the nodes.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

The slot number shown is the slot number in use on stacking enabled nodes. If the node does not have stacking enabled, a hyphen character (-) is shown instead of a number.

The Port and Node MAC Address field values in the command display identify a particular stacking port. Each node MAC address appears twice in two consecutive rows in the output because each node has two stacking ports. On all platforms, the ports are labeled with the values 1 or 2. The order in which stacking ports appear in the display is the order in which they are physically connected.

The Select field indicates whether the stacking port is using a native stacking port or an alternate 10Gbps Ethernet port. If a number appears in this column, it represents the port number printed on the switch for a 10 Gbps Ethernet port. For more information, see the description for the [configure stacking-support stack-ports](#) command.

The Port State field for each port shows one of the following states:

- Link Down – port is not receiving a signal.
- No Neighbor – the port is receiving a signal but it is not identifying a stack neighbor.
- Overflow – 17 nodes (or more) are physically connected to this port.

- Inhibited – When you connected the link, active topologies were detected on both sides, and at least one slot number was duplicated. The stack merge is blocked.
- Operational – the port is operational in the stack. This is a necessary but insufficient condition for the port to be used for control path or user data. For example, a node with stacking Failed state may still show its port states as Operational.

The Flags field contains the following flag definitions:

- C - The control path is active on this port. Note that the user data path over the stack links follows the control path.
- B - The port is blocked from transmitting traffic that is to be flooded to multiple non-stacking ports. This flag is only set in an active ring topology on two adjacent ports. In the example below, the active topology is a daisy chain, so no ports are blocked.

Identified with the asterisk, the current node is the one on which the show stacking command was executed. The stack topology is shown in a particular order. In a ring topology, the current node is always the first node, the next node is the node connected to the port 2 of the first node, and the last node is the node connected to the port 1 of the current node. In a daisy chain, the order shown depends on the connection of the node on which the command executes:

- The first node is the one at the far end of the daisy-chain connected to the current node port 1.
- The last node is the one at the far end of the daisy-chain connected to the current node port 2.
- The previous node is the one at the near end of the daisy-chain connected to the current node port 1.
- The next node is the one at the near end of the daisy-chain connected to the current node port 2.
- If there is no node connected to the current node port 1, the current node is the first node.
- If there is no node connected to the current node port 2, the current node is the last node.

The port speed is the unidirectional speed of the port.



Note

Some VIM names include speed ratings which are 4 times the unidirectional stacking port speed. For example, the actual stacking port speed for VIM1-SummitStack512 is 128 Gbps. The 512 Gbps rating for the VIM is the unidirectional rate X 2 (bidirectional) X 2 (ports).

Example

The following example shows the command output for a stack that is operating in a ring and uses both native and alternate stack ports:

```
Slot-1 Stack.9 # show stacking stack-ports
Stack Topology is a Ring
Slot Port Select Node MAC Address Port State Flags Speed
-----
*1 1 23 00:04:96:26:6b:ec Operational C- 10G
*1 2 Native 00:04:96:26:6b:ec Operational CB 64G
2 1 Native 00:04:96:18:7d:e8 Operational CB 64G
2 2 24 00:04:96:18:7d:e8 Operational C- 10G
3 1 23 00:04:96:27:c5:12 Operational C- 10G
3 2 Native 00:04:96:27:c5:12 Operational C- 64G
4 1 Native 00:04:96:26:6b:34 Operational C- 64G
4 2 24 00:04:96:26:6b:34 Operational C- 10G
```

```
* - Indicates this node
Flags: (C) Control path is active, (B) Port is Blocked
Slot-1 Stack.10 #
```

The following example shows the command output for stacks that use the 512 Gbps stacking ports:

```
Slot-1 Stack.2 # show stacking stack-ports
Stack Topology is a Ring
Slot Port Select Node MAC Address Port State Flags Speed
-----
*1 1 Native 00:04:96:35:8b:a5 Operational C- 128G
*1 2 Native 00:04:96:35:8b:a5 Operational C- 128G
2 2 Native 00:04:96:35:a8:b0 Operational C- 128G
2 1 Native 00:04:96:35:a8:b0 Operational C- 128G
```



Note

Although the VIM1-SummitStack512 option card has four physical ports, the physical ports are grouped into two pairs, forming two logical ports. The show stacking stack-ports command displays the status of the logical ports.

History

This command was first available in ExtremeXOS 12.0.

The Select column was added in ExtremeXOS 12.5.

Platform Availability

This command is available with all licenses and platforms that support the SummitStack feature. For information about which licenses and platforms support the SummitStack feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

show stacking-support

```
show stacking-support
```

Description

This command displays the configured and current states of configuration options configured on the local node with the **stacking-support** keyword.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

The display parameters are described in the following table:

Display Item	Description
Stack Port column	Displays rows for Stack Port 1 and Stack Port 2.
Native column	Indicates whether the switch has native stack ports. A Yes entry indicates that the switch has native stacking ports. A No entry indicates that no native stacking ports are present. An asterisk indicates that the native stack port is selected.
Alternate column	Displays the port numbers for data ports that can operate as alternate stack ports. A No entry indicates that there are no data ports that can operate as alternate stacking ports. An asterisk indicates that the data port number is selected as an alternate stack port.
Configured column	Indicates the configured state for stack ports 1 and 2, which can be Native or Alternate. This column also indicates the configured state of stacking-support option. For platform configurations with dual-purpose hardware (that supports stack ports or data ports), this column displays either Enabled (stack ports enabled) or Disabled (stack ports disabled, data ports active). For platform configurations without dual-purpose hardware, this column displays N/A, which indicates that stacking-support option cannot be disabled on this switch. This is the configuration that becomes active the next time the switch boots if the stacking-support option is enabled. For more information, see the command descriptions for the enable stacking-support and disable stacking-support commands.
Current column	Indicates the selection that is currently in effect for stack ports, which can be Native, Alternate, or N/A. N/A indicates that the port selection is not applicable to this switch hardware configuration. The column also indicates the current operating state of the stacking-support option, which can be Enabled, Disabled, or N/A. An N/A entry indicates that no option card is present. This is the configuration that is active now.

Example

The following example shows the stack port selection and stacking-support option configuration after the unconfigure stacking-support command has been executed and before a subsequent reboot has been initiated:

```
# show stacking-support
Stacking Support Settings
Stack   Available Ports
Port    Native  Alternate  Configured  Current
-----
1       Yes *   23        Native      Native
2       Yes    24 *      Native      Alternate
stacking-support:      Disabled    Enabled
auto-discovery:        Enabled     Disabled
Flags: * - Current stack port selection
NOTE: This node must be rebooted before the configured settings will
take effect.
```

The following example shows that the stacking-support option is disabled and will remain disabled when the switch reboots:

```
# show stacking-support
Stacking Support Settings
Stack   Available Ports
Port    Native  Alternate  Configured  Current
-----  -
1       No     S3         Native      N/A
2       No     S4         Native      N/A
stacking-support:      Disabled  Disabled
auto-discovery:        Enabled   Enabled
Flags: * - Current stack port selection
```

The following example shows the configured and current pseudo native stack port type:

```
# show stacking-support

Stack   Available Ports
Port    Native  Alternate  Configured  Current
-----  -
1       Yes *   47         Native (V160) Native (V80)
2       Yes *   48         Native      Native (V320)
stacking-support:      Enabled   Enabled
auto-discovery:        Enabled   Enabled

Flags: * - Current stack port selection
```

History

This command was first available in ExtremeXOS 12.5.

Ability to show configured and current pseudo native stack port type (V160, V320, V400) was added in ExtremeXOS 30.1.

Stacking auto-discovery information was added in ExtremeXOS 30.3.

Platform Availability

This command is available on all platforms that support alternate stack port selection or permit disabling of the stacking-support option.

show stpd ports blocked-ports

```
show stpd stpd_name ports {port_list} {blocked-ports}
```

Description

Displays all blocked ports in Spanning Tree.

Syntax Description

<i>stpd_name</i>	Specifies an <i>STPD</i> name on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.
blocked-ports	Displays only blocked ports.

Default

N/A

Usage Guidelines

None.

Example

The following example displays the blocked Spanning Tree ports:

```
show stpd r1 blocked-ports

Port      VLAN
-----
1:1      v123456789012345678901234567890, v1234567890123456789012345678901,
        v1234567890123456789012345678901

Number of Blocked Ports : 3
```

History

This command was first available in ExtremeXOS 15.7.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show stpd ports counters

```
show stpd stpd_name ports [port_list] counters
```

Description

Displays all counters for Spanning Tree.

Syntax Description

<i>stpd_name</i>	Specifies an <i>STPD</i> name on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.
counters	Displays counters for the specified port.

Default

N/A

Usage Guidelines

None.

Example

The following example displays the counters for port 10:10 in *STP* domain r1:

```
show stpd r1 ports 10:10 counters

Port 10:10
-----
Disputed BPDU           : 0
Invalid BPDU            : 0
Message Expiration      : 0
STP BPDU Rx             : 0
STP BPDU Tx             : 0
STP TCN BPDU Rx         : 0
STP TCN BPDU Tx         : 0
STP TC BPDU Rx          : 0
STP TC BPDU Tx          : 0
RST BPDU Rx             : 0
RST BPDU Tx             : 0
RST TC BPDU Rx          : 0
RST TC BPDU Tx          : 0
MST BPDU Rx             : 1287
MST BPDU Tx             : 4
MST CIST TC BPDU Rx     : 2
MST CIST TC BPDU Tx     : 2
Forward Transitions Count : 1
```

History

This command was first available in ExtremeXOS 15.7.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show stpd ports non-forwarding-reason

```
show stpd stpd_name ports {port_list} { non-forwarding-reasons }
```

Description

Displays the reasons for placing a port in a non-forwarding state.

Syntax Description

<i>stpd_name</i>	Specifies an <i>STPD</i> name on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.
non-forwarding-reasons	Displays only reasons for placing a port in a non-forwarding state.

Default

N/A

Usage Guidelines

Ports can be placed in the non-forwarding state for the following reasons:

- Placed in listening state because it is in dispute.
- Not been placed in a non-forwarding state due to any exceptional condition.
- Placed in listening or blocking state by the Loop Protect feature.
- Placed in listening state by the Loop Protect feature, but this condition may be normal because the link partner does not support Loop Protect.
- Placed in blocking state because a loopback condition has been detected.
- Unknown non-forwarding reason.

Example

The following example displays the non-forwarding reasons on port 1:10 in *STP* domain r1:

```
show stpd r1 ports 1:10 non-forwarding-reason
```

Port	Reason
1:1	Placed in listening state because it is in dispute.
1:2	Not been placed in a non-forwarding state due to any exceptional condition.
1:3	Placed in listening or blocking state by the Loop Protect feature.
1:4	Placed in listening state by the Loop Protect feature, but this condition may be normal because the link partner does not support Loop Protect.
1:5	Placed in blocking state because a loopback condition has been detected.
1:6	Unknown non-forwarding reason.

History

This command was first available in ExtremeXOS 15.7.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show stpd

```
show stpd {stpd_name | detail}
```

Description

Displays *STPD* settings on the switch.

Syntax Description

<i>stpd_name</i>	Specifies an STPD on the switch.
detail	Specifies that STPD settings should be shown for each STPD.

Default

N/A.

Usage Guidelines

If you specify the command without any options, the following STPD information appears:

- Name—The name of the STPD.
- Tag—The StpdID of the domain, if configured.
- Flags—The following flags communicate information about the current state of the STPD:
 - (C) Topology Change—A network topology change has occurred in the network.
 - (D) Disable—The STPD is disabled.
 - (E) Enable—The STPD is enabled.
 - (R) Rapid Root Failover—The STPD has been configured for rapid root failover.
 - (T) Topology Change Detected—The STPD has detected a change in the network topology.
 - (M) *MSTP* CIST—The STPD has been configured for MSTP, and the STPD is the common and internal spanning tree.
 - (I) MSTP *MSTI*—The STPD has been configured for MSTP, and the STPD is a multiple instance spanning tree.
- Ports—The number of ports that are part of the STPD.
- Bridge ID—The MAC addresses of the switch.
- Designated Root—The MAC address of the switch that is the designated root bridge.
- Rt Port—The root port.
- Rt Cost—The path cost to the root port.

- Total Number of STPDs—The total number of STPDs configured on the switch.
- STP Flush Method—The method used to flush the FDB during a topology change.

If you have an MSTP region and associated spanning trees configured on the switch, the command also displays the following global MSTP information:

- MSTP Region—The name of the MSTP region configured on the switch.
- Format Identifier—The number used by BPDUs to communicate within an MSTP region.
- Revision Level—This number is reserved for future use.
- Digest—The MD5 digest value.
- Common and Internal Spanning Tree (CIST)—The name of the CIST that controls the connectivity of interconnecting MSTP regions.
- Total number of MST Instances (MSTI)—The number of MSTIs running in the MSTP region.

If you use the `show stpd` command and specify the name of an STPD, in addition to the data previously described, the command displays more detailed information about the STPD. If you specify the `detail` option, the switch displays the same type of information for all of the STPDs configured on the switch.

The additional output includes the following:

- STPD mode of operation.
- Autobind mode.
- Active VLANs.
- Bridge priority mode.
- Timer information.
- Topology change information.

If you have MSTP configured, the command also displays the following information:

- Bridge role.
- CIST root.
- CIST regional root.
- MSTI instances.
- Master port (Displayed only on MSTI STPDs).

If your STPD has the same name as another component, for example a VLAN, we recommend that you specify the identifying keyword as well as the name. If you do not specify the `stpd` keyword, an error message similar to the following is displayed:

```
%% Ambiguous command: "show Test"
```

In this example, to view the settings of the STPD `Test`, enter `show stpd Test`.

If your STPD has a name unique only to that STPD, the keyword `stpd` is optional.

Example

The following command displays the STPD settings on a switch that has MSTP configured:

```
# show stpd
```

The following is sample output from this command:

```
MSTP Global Configuration:
MSTP Region Name       : 00049651acd7
MSTP Format Identifier  : 0
MSTP Revision Level    : 3
MSTP Digest            : ac:36:17:7f:50:28:3c:d4:b8:38:21:d8:ab:26:de:62
Common and Internal Spanning Tree (CIST) : ----
Total Number of MST Instances (MSTI)      : 0

Name      Tag  Flags  Ports Bridge ID      Designated Root  Rt Port Rt Cost
s0        0000 D----- 0 800000049651acd7 0000000000000000 ----- 0

Total number of STPDs: 1          STP Flush Method: VLAN and Port
STP Filter Method: System-wide    STP BPDU Forwarding: On
STP Multicast Send IGMP or MLD Query: On

Flags: (C) Topology Change, (D) Disable, (E) Enable, (R) Rapid Root Failover
       (T) Topology Change Detected, (M) MSTP CIST, (I) MSTP MSTI
```

The following command displays STPD settings on an STPD named Backbone_st:

```
show stpd backbone_st
```

The following is sample output from this command:

```
Stpd: backbone_st Stp: ENABLED      Number of Ports: 51
Description: this is backbone_st domain
Rapid Root Failover: Disabled
Operational Mode: 802.1W Default Binding Mode: 802.1D
802.1Q Tag: (none)
Ports: 1:1,1:2,2:1,2:2,3:1,3:2,4:1,4:2,5:1,5:2,
5:3,5:4,5:5,5:6,5:7,5:8,5:9,5:10,5:11,5:12,
5:13,5:14,5:15,5:16,5:17,5:18,5:19,5:20,5:21,5:22,
5:23,5:24,5:25,5:26,5:27,5:28,5:29,5:30,5:31,5:32,
5:33,5:34,5:35,5:36,5:37,5:38,5:39,5:40,5:41,5:42,
5:43
Participating Vlan: Default
Auto-bind Vlan: Default
Bridge Priority: 5000
c
BridgeID: 13:88:00:01:30:f4:06:80
Designated root:      0a:be:00:01:30:28:b7:00
RootPathCost: 19      Root Port: 28
MaxAge: 20s           HelloTime: 2s           ForwardDelay: 15s
CfgBrMaxAge: 20s      CfgBrHelloTime: 2s          CfgBrForwardDelay: 15s
Topology Change Time: 35s Hold time: 1s
Topology Change Detected: FALSE Topology Change: FALSE
Number of Topology Changes: 7
Time Since Last Topology Change: 4967s
```

The following command displays detailed output for STP domain s0.

```
show stpd s0 detail
```

The following is sample detailed output.

```
Stpd: s0              Stp: DISABLED      Number of Ports: 0
Rapid Root Failover: Disabled
Operational Mode: 802.1D          Default Binding Mode: 802.1D
802.1Q Tag: (none)
Ports: (none)
```

```

Participating Vlans: v1,v11,v21
Auto-bind Vlans: Default
Bridge Priority: 32768
Operational Bridge Priority : 28672
Bridge Priority Mode: 802.1d
BridgeID: 70:00:00:04:96:82:6a:29
Designated root: 00:00:00:00:00:00:00:00
RootPathCost: 0 Root Port : ----
MaxAge : 0s HelloTime : 0s ForwardDelay : 0s
CfgBrMaxAge : 20s CfgBrHelloTime: 2s CfgBrForwardDelay: 15s
Topology Change Time : 35s Hold time : 1s
Topology Change Detected : FALSE Topology Change : FALSE
Number of Topology Changes : 0
Time Since Last Topology Change: 0s
Topology Change initiated locally on Port 1:21
Topology Change last received on Port 1:21 from 00:04:96:26:6c:89
Backup Root : On Backup Root Activated : TRUE
Backup Root Trap : On New Root Trap : On
BPDU Restrict Trap : On Topology Change Trap : On
Dispute Threshold : None Loop Protect Threshold : None
Loop Protect Event Window : 180s Loop Protect Trap : On
Tx Hold Count : 6
Participating VLANs :
VLAN Tag Number of Ports
v1 100 13
1:1 (F), 1:2 (B), 1:3 (L), 1:4 (B), 1:5 (D), 1:6 (D), 1:7 (D), 1:8 (D), 1:9 (B), 1:10 (L),
1:11 (F), 1:12 (F), 1:13 (F)
v11 1001 9
1:23 (F), 1:24 (B), 1:25 (L), 1:26 (B), 1:6 (D), 1:7 (D), 1:8 (D), 1:9 (B), 1:10 (L),
v21 2001 3
1:23 (F), 1:24 (B), 1:25 (L)
Flags: B-Blocking, D-Disabled, F-Forwarding, I-Listening, L-Learning.

```

The following is sample output for an STPD configured as the CIST (the output is similar for an STPD configured as an MSTI):

```

Stpd: s0 Stp: DISABLED Number of Ports: 0
Description: this is s0 domain
Rapid Root Failover: Disabled
Operational Mode: MSTP Default Binding Mode: 802.1d
MSTP Instance :CIST CIST : s0
802.1Q Tag: (none) Ports:
(none) Participating
Vlan Count: 1
Auto-bind Vlans Count: 1
Bridge Priority: 32768 BridgeID:
80:00:00:10:30:f9:9d:c0Bridge
Role : CIST Regional Root
CIST Root 80:00:00:10:30:f9:9d:c0CIST
Regional Root: 80:00:00:10:30:f9:9d:c0
Designated root: 00:00:00:00:00:00:00:00 RootPathCost:
0 External RootPathCost: 0 Root Port: ----
MaxAge:0sHelloTime:0sForwardDelay:0s
CfgBrMaxAge:20sCfgBrHelloTime:2sCfgBrForwardDelay: 15s MaxHopCount: 20
CfgBrMaxHopCount : 20
Topology Change Time: 35s Hold time:
1s Topology Change Detected: FALSE Topology Change: FALSE
Number of Topology Changes: 0 Time
Since Last Topology Change: 0s
Participating Vlans :
(none) Auto-bind Vlans : Default

```

History

This command was first available in ExtremeXOS 10.1.

Information about MSTP was added in ExtremeXOS 11.4.

Description was added in ExtremeXOS 12.4.4.

MSTP Digest, topology change information, and trap information were added in ExtremeXOS 15.7.1.

Bridge priority mode information added in ExtremeXOS 22.1.

Send IGMP/MLD query suppression status added in ExtremeXOS 22.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show stpd ports

```
show {stpd} stpd_name ports {[detail | port_list {detail}]}
```

Description

Displays the *STP* state of a port.

Syntax Description

<i>stpd_name</i>	Specifies an <i>STPD</i> name.
<i>port_list</i>	Specifies one or more ports or slots and ports.
detail	Specifies more detailed information about one or more ports of the STPD.

Default

N/A.

Usage Guidelines

This command displays the following:

- STPD port configuration.
- STPD port encapsulation mode.
- STPD path cost.
- STPD priority.
- STPD state (root bridge, and so on).
- Port role (root designated, alternate and so on).
- STPD port state (forwarding, blocking, and so on).

- Configured port link type.
- Operational port link type.
- Edge port settings (inconsistent behavior, edge safeguard setting).
- Restricted role (enabled, disabled).
- MSTP port role (internal or boundary).
- Active port role.

To display more detailed information for one or more ports in the specified STPD, including participating VLANs, specify the **detail** option.

If you have MSTP configured and specify the detail option, this command displays additional information:

- MSTP internal path cost.
- MSTP timers.

If your STPD has the same name as another component, for example a VLAN, Extreme Networks recommends that you specify the identifying keyword as well as the name. If you do not specify the stpd keyword, an error message similar to the following is displayed:

```
%% Ambiguous command: "show Test ports"
```

In this example, to view all of the port settings of STPD Test, enter show stpd Test ports.

If your STPD has a name unique only to that STPD, the keyword stpd is optional.

Example

The following command displays the state of ports 1, 2, and 4 on an STPD named s1:

```
show stpd s1 ports
```

The following is sample output from this command:

```
Port   Mode   State      Cost  Flags      Priority Port ID Designated Bridge
1      EMISTP DISABLED 200000 e?pp-w---t 128      8001    00:00:00:00:00:00:00
2      EMISTP DISABLED 200000 e?pp-w---- 128      8002    00:00:00:00:00:00:00
4      EMISTP DISABLED 200000 e?pp-w---- 128      8004    00:00:00:00:00:00:00
Total Ports: 3
----- Flags: -----
1:          e=Enable, d=Disable
2: (Port role) R=Root, D=Designated, A=Alternate, B=Backup, M=Master
3: (Config type) b=broadcast, p=point-to-point, e=edge, a=auto
4: (Oper. type) b=broadcast, p=point-to-point, e=edge
5:          p=proposing, a=agree
6: (partner mode) d = 802.1d, w = 802.1w, m = mstp
7:          i = edgeport inconsistency
8:          S = edgeport safe guard active
s = edgeport safe guard configured but inactive
8:          G = edgeport safe guard bpdu restrict active in 802.1w and mstp
g = edgeport safe guard bpdu restrict active in 802.1d
9:          B = Boundary, I = Internal
10:         r = Restricted Role, t = active Role
```

The following command displays the detailed information for port 2 in STPD s0:

```
show stpd s0 ports 2 detail
```

The following is sample output from this command:

```
show stpd s0 ports 2 detail
Stpd: s0          Port: 2 PortId: 8002    Stp: ENABLED    Path Cost: 4
Port Mode        : 802.1D                Port Role       : ----
Port State       : FORWARDING             Topology Change Ack: FALSE
Port Priority     : 16
Designated Root  : 00:00:00:00:00:00:00:00 Designated Cost: 0
Designated Bridge : 00:00:00:00:00:00:00:00 Designated Port Id: 0
Partner STP version : Dot1d
Restricted Role     : Disabled
Active Role        : Disabled
Edge Port Safe Guard : Disabled
BPDU Restrict      : Disabled
Restricted TCN      : Off
Loop Protect       : Off
Loop Protect Partner : Incapable
Operational Edge   : FALSE
Auto Edge          : On
Reflection BPDU    : On
Participating Vlans: Default
```

The following is sample output from this command:

```
Port  Mode  State      Cost  Flags      Priority Port ID Designated Bridge
9      EMISTP FORWARDING 20000 eDeepw-G-- 128      8009    80:00:00:04:96:1f:a8:48
Total Ports: 1
----- Flags: -----
1:          e=Enable, d=Disable
2: (Port role) R=Root, D=Designated, A=Alternate, B=Backup, M=Master
3: (Config type) b=broadcast, p=point-to-point, e=edge, a=auto
4: (Oper. type) b=broadcast, p=point-to-point, e=edge
5:          p=proposing, a=agree
6: (partner mode) d = 802.1d, w = 802.1w, m = mstp
7:          i = edgeport inconsistency
8:          S = edgeport safe guard active
s = edgeport safe guard configured but inactive
G = edgeport safe guard bpdu restrict active
g = edgeport safe guard bpdu restrict configured but inactive only dot1w, mstp
9:          B = Boundary, I = Internal
10:         r = Restricted Role, t = active role
```

History

This command was first available in ExtremeXOS 10.1.

Information about MSTP was added in ExtremeXOS 11.4.

Information about BPDU Restrict was added in ExtremeXOS 12.4.

Information about active role was added in ExtremeXOS 12.5.

Information about reflection BPDU status was added in ExtremeXOS 22.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show switch

```
show switch {detail}
```

Description

Displays the current switch information.

On a SummitStack, this command displays the Master and Backup node information if executed on the Master, and displays the current node and the Master node information if executed on any other node.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

The show switch command displays:

- sysName, sysLocation, sysContact
- MAC address
- System type
- System health check
- Recovery mode
- Watchdog state
- Secure Boot (trusted delivery) status (5420, 5520, and 5720 switches only):
 - UBOOT Image Verified



Note

If the UBOOT image verification fails during boot-up, the switch halts and enters the Recovery Boot Loader.

- Current date, time, system boot time, and time zone configuration
- Any scheduled reboot information
- System up time
- Master and Backup information (available only on SummitStack)
- Current state (available only on stand-alone switches)
 - OPERATIONAL.
 - OPERATIONAL (OverHeat).
 - FAILED.

- Software image information (primary/secondary image and version)
- Configuration information (primary/secondary configuration and version)
- If locally administered MAC address generation is enabled (see [enable switch locally-administered-address](#) on page 2336)
- Software version information

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

On a SummitStack, the System UpTime may be useful when manually resolving the dual master situation. For more information, see *Eliminating a Dual Master Situation Manually* section in the [Switch Engine 32.3 User Guide](#).

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

On a stack or Extended Edge Switching topology, the following additional information is available:

- System Type
- System UpTime
- Details of master and backup, or current node and master, including any loss of synchronicity with the primary, and which process is out of sync

Example

Output from this command on the standalone ExtremeSwitching series switch looks similar to the following:

```
SysName:          SummitX
SysLocation:
SysContact:       support@extremenetworks.com, +1 888 257 3000
System MAC:       00:04:96:26:6B:EC
System Type:      SummitX
SysHealth check:  Enabled (Normal)
Recovery Mode:    All
System Watchdog:  Enabled
Trusted Delivery: UBOOT Image Verified

Current Time:     Wed Apr 25 21:17:18 2012
Timezone:         [Auto DST Disabled] GMT Offset: 0 minutes, name is UTC.
Boot Time:        Wed Apr 25 21:13:54 2012
Boot Count:       951
Next Reboot:      None scheduled
System UpTime:    3 minutes 24 seconds
Current State:    OPERATIONAL
Image Selected:   secondary
Image Booted:     secondary
Primary ver:      12.0.0.4
Secondary ver:    12.0.0.4Config Selected:  primary.cfg
Config Booted:    primary.cfg
Config Automatic: primary.cfg
primary.cfg       Created by ExtremeXOS version 22.2.0.16
                  344404 bytes saved on Tue Jan 17 11:17:56 2017
                  Auto-saved every 2 minutes.
                  Next periodic save on Tue Jan 17 14:45:33 2017
LAA MAC:          Locally Administered MAC Address Enabled
Port Name Pattern: eths\sp\p (After Reboot: slot\sport\p)
show version
```



```

Slot-1      : 800908-00-01 1847F-10076 Rev 01 IMG: 30.3.0.453
VIM5-4XE-1  : 800910-00-01 1845F-10106 Rev 01
XN-SSD-001-120-1: 800954-00-02 1917F-10030 Rev 02

Image       : ExtremeXOS version 30.3.0.453 by release-manager
              on Thu Jul 11 11:52:10 EDT 2019
Diagnostics :
Certified Version : EXOS Linux 4.14.123, FIPS fips-ecp-2.0.16

Build Tools Version : exos-x32-sdk-2.5.3.1.0

```

The `show switch detail` command displays the same information shown above. Output from this command on a stack looks similar to the following:

```

SysName:      Stack
SysLocation:
SysContact:   support@extremenetworks.com, +1 888 257 3000
System MAC:   02:04:96:27:B7:41
System Type:  X (Stack)
SysHealth check: Enabled (Normal)
Recovery Mode: All
System Watchdog: Enabled
Current Time: Tue Jan 30 14:22:41 2007
Timezone:     [Auto DST Disabled] GMT Offset: 0 minutes, name is UTC.
Boot Time:    Mon Jan 29 21:51:38 2007
Boot Count:   317
Next Reboot:  None scheduled
System UpTime: 16 hours 31 minutes 3 seconds
Slot:         Slot-4 *                               Slot-5
-----
Current State: MASTER                                BACKUP (In Sync)
Image Selected: secondary                            secondary
Image Booted:   secondary                            secondary
Primary ver:    12.0.0.10                             12.0.0.10
Secondary ver:  12.0.0.13                             12.0.0.13
Config Selected: primary.cfg
Config Booted:  primary.cfg
primary.cfg     Created by ExtremeXOS version 12.0.0.10
139108 bytes saved on Fri Jan 26 22:56:40 2007
LAA MAC:        Locally Administered MAC Address Disabled
Port Name Pattern: eths\sp\p (After Reboot: eths\sp\p)

```

Output for a Extended Edge Switching topology looks similar to the following:

```

Slot-1 VPEX Stack.12 # sh switch

SysName:      Stack
SysLocation:
SysContact:   https://www.extremenetworks.com/support/
System MAC:   02:04:96:9B:B7:87
System Type:  VPEX X670G2-48x-4q (Stack)

SysHealth check: Enabled (Normal)
Recovery Mode:   All
System Watchdog: Enabled

Current Time:    Wed May 6 23:54:50 2020
Timezone:        [Auto DST Disabled] GMT Offset: 0 minutes, name is UTC.
Boot Time:       Wed May 6 23:15:57 2020
Boot Count:      2405
Next Reboot:     None scheduled
System UpTime:   38 minutes 53 seconds

```

Slot:	Slot-1 *	Slot-2
	-----	-----
Current State:	MASTER	BACKUP (vpex Not Synced)
Image Selected:	secondary	secondary
Image Booted:	secondary	secondary
Primary ver:	30.7.0.102	30.7.0.102
Secondary ver:	31.1.0.108	31.1.0.108
Config Selected:	primary.cfg	
Config Booted:	primary.cfg	
Config Automatic:	NONE (Disabled)	
primary.cfg	Created by ExtremeXOS version 31.1.0.74 951802 bytes saved on Mon Apr 20 19:34:02 2020	

History

This command was first available in ExtremeXOS 10.1.

This command was updated to support stacking in ExtremeXOS 12.0 and the System Type was added to the output from this version.

Automatic configuration save information was added in ExtremeXOS 22.2.

Locally administered per-port MAC address information was added in ExtremeXOS 22.3.

System port name pattern information added in ExtremeXOS 30.1.

Solid State Storage Device SSD-120 information was added in ExtremeXOS 30.3.

Secure Boot (trusted delivery) status information and additional process out-of-sync information for stacks and Extended Edge Switching was added for ExtremeSwitching 5520 series switches in ExtremeXOS 31.1, for ExtremeSwitching 5420 series switches in ExtremeXOS 31.3, and for ExtremeSwitching 5720 series switches in Switch Engine 32.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show switch bluetooth

```
show switch bluetooth [statistics | inventory]
```

Description

Displays Bluetooth device-related information.

Syntax Description

switch	Designates showing information for the switch.
bluetooth	Designates showing information about Bluetooth devices.

statistics	Shows Bluetooth statistics.
inventory	Shows Bluetooth inventory (model number, serial number, etc.).

Default

N/A.

Usage Guidelines

Using this command without the **statistics** shows:

- If Bluetooth is enabled on the switch.
- Switch MAC address.
- Switch name.
- Is Bluetooth discovery enabled.
- Is Bluetooth pairing enabled.
- List of Bluetooth devices:
 - Device MAC address.
 - Device name.
 - Device type (PC, phone, etc.)
 - Is the device paired.
 - Is the device trusted.
 - Is the device blocked.
 - Is the device paired.

To enable Bluetooth capabilities, use the enable **switch bluetooth {discovery | pairing }** command.

To disable Bluetooth capabilities, use the disable **switch bluetooth {discovery | pairing }** command.

Example

The following example shows Bluetooth information for a switch:

```
# show switch bluetooth
Bluetooth      : Enabled
  Address      : 00:04:96:9a:46:48
  Name         : X450G2-24p-G4
  Discovery    : Enabled
  Pairing      : Enabled
```

Bluetooth Address	Device Name	Device Type	Paired	Trusted	Blocked	Connected
00:00:00:00:00:02	Device1	PC	Yes	Yes	No	Yes
00:00:00:00:00:03	Device2	Phone	Yes	No	No	No

The following example shows Bluetooth statistics:

```
# show switch bluetooth statistics
Bluetooth          Rx      Tx      Rx      Err      Tx      Rx      Tx
Address           Bytes   Bytes  Events    Cmd    ACL    ACL
-----
00:04:96:9a:46:48    50    100     10      10      10     15     15
```

The following example shows Bluetooth inventory:

```
# show switch bluetooth inventory
Bluetooth
Address      Manufacturer  Model      Model #    Serial
#            Description
-----
5C:F3:70:8B:E4:11 Broadcom_Corp  BCM20702A0  21e8
Broadcom_Corp_BCM20702A0_5CF3708 usb_device
```

History

This command was first available in ExtremeXOS 30.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show switch management

```
show switch management
```

Description

Shows management information.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Example

The following example shows switch management information:

```
# show switch management

CLI idle timeout      : Enabled (20 minutes)
```

```

CLI max number of login attempts : 3
CLI max number of sessions      : 8
CLI paging                      : Enabled
CLI space-completion            : Disabled (this session only)
CLI configuration logging        : Disabled (without expansion)
CLI journal size                 : 100
CLI password prompting only     : Disabled
CLI display moved-keywords      : Hidden
CLI moved-keywords hidden release: 31.1
CLI RADIUS cmd authorize tokens : 2
CLI scripting                   : Disabled (this session only)
CLI scripting error mode         : Ignore-Error (this session only)
CLI script search path           : ".:usr/local/cfg" (this session only)
CLI persistent mode              : Persistent (this session only)
CLI prompting                   : Enabled (this session only)
CLI screen size                  : 24 Lines 80 Columns (this session only)
CLI refresh                     : Enabled
CLI history expansion            : Disabled
Image integrity checking         : On (Valid)
Current system port notation    : slot:port
Telnet access                    : Enabled (tcp port 23 vr all)
                                : Access Profile : not set
SSH access                      : Disabled (Key invalid, tcp port 22 vr all)
                                : Secure-Mode   : Off
                                : Access Profile : not set
SSH2 idle time                  : 60 minutes
SSH2 rekey interval             : 4096 MB and no time limit
Web access                      : Enabled (tcp port 80)
                                : Access Profile : not set
Total Read Only Communities    : 0
Total Read Write Communities   : 0
RMON                           : Disabled
SNMP access                    : Enabled
                                : Access Profile : not set
SNMP Notifications             : Enabled
SNMP Notification Receivers    : None
SNMP stats:      InPkts 0      OutPkts 0      Errors 0      AuthErrors 0
                  Gets   0      GetNexts 0      Sets   0      Drops     0
SNMP traps:      Sent   0      AuthTraps Enabled
SNMP inform:     Sent   0      Retries  0      Failed  0

```

History

This command was first available in ExtremeXOS 10.1.

Image integrity check information was added in ExtremeXOS 31.1.

Script search path was added in ExtremeXOS 31.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show switch mounts

```
show switch mounts
```

Description

Displays whether a USB 2.0 storage device is present on the switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

This command shows whether a removable USB storage device is present on the switch:

```
# show switch mounts
```

The following output appears.

If you do not have a removable storage device installed:

```
Memory storage is not present.
```

If only one USB is supported and present:

```
Memory storage is present.
```

When two USBs are supported and present:

```
Memory storage is present in USB-1.  
Memory storage is present in USB-2.
```

History

This command was first available in ExtremeXOS 11.0.

Support for USB 2.0 storage devices was added in ExtremeXOS 12.5.3.

The **memorycard** keyword was deprecated and replaced with the keyword **mounts** in ExtremeXOS 30.7.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show switch usb

```
show switch usb
```

Description

Shows the status (enabled or disabled) of the switch's USB port.

Syntax Description

usb	Specifies USB port on switch.
------------	-------------------------------

Default

N/A.

Usage Guidelines

The enable/disable switch usb commands require a reboot to take effect. This is reflected in the output of this show command:

```
# show switch usb
USB port: Enabled
# disable switch usb
This setting will take effect at the next system reboot.
# show switch usb
USB port: Enabled (Disabled after the next reboot)Or in the same output, you can
add:"After disabling usb using 'disable switch usb' output will be like this:# show
switch usbUSB port: Enabled (Disabled after the next reboot)
```

Example

The following example shows the switch's USB port status:

```
# show switch usb
USB port: Enabled
```

History

This command was first available in ExtremeXOS 22.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show system

```
show system
```

Description

This command displays the aggregated output of the following commands:

- `show switch`
- `show version`
- `show temperature`
- `show power`
- `show fans`
- `show odometers`

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to view the collected output of individual system commands.

Example

```
# show system

show switch
SysName:          5520-24W-EXOS
SysLocation:
SysContact:       https://www.extremenetworks.com/support/
System MAC:       00:04:96:EF:2C:00
System Type:      5520-24W-EXOS

SysHealth check:  Enabled (Normal)
Recovery Mode:    All
System Watchdog:  Enabled
Trusted Delivery: UBOOT Image Verified

Current Time:     Tue Sep  1 14:51:40 2020
Timezone:         [Auto DST Disabled] GMT Offset: 0 minutes, name is UTC.
Boot Time:        Tue Aug 11 18:18:57 2020
Boot Count:       105
Next Reboot:      None scheduled
System UpTime:    20 days 20 hours 32 minutes 43 seconds

Current State:    OPERATIONAL
Image Selected:   secondary
Image Booted:     secondary
Primary ver:      31.1.0.398
Secondary ver:    31.1.0.485

Config Selected:  primary.cfg
Config Booted:    primary.cfg
Config Automatic: NONE (Disabled)
```



```

primary.cfg          Created by ExtremeXOS version 31.1.1.142
                    354129 bytes saved on Mon Jul 27 13:57:50 2020

LAA MAC:             Locally Administered MAC Address Disabled

show version
Switch               : 800992-01-03 2011G-00108 Rev 03 BootROM: 2.2.1.3   IMG: 31.1.0.485

Image   : ExtremeXOS version 31.1.0.485 by release-manager
        on Thu Aug 6 07:07:23 EDT 2020
BootROM : Default 2.2.1.3 Alternate 2.2.1.3
Diagnostics :
Certified Version : EXOS Linux 4.14.170, Extreme Networks FIPS Object Module 2.0.16a
Build Tools Version : exos-arm64-sdk-3.0.3.1.0

show temperature
Field Replaceable Units          Temp (C)   Status   Min
Normal   Max
-----
Switch   : 5520-24W-EXOS          43.00    Normal   0
10-100   110

show power
PowerSupply 1 information:
State           : Powered On
PartInfo        : Internal PSU-1 1943A-33011 800592-00-15
Revision        : 15
Input           : 208.25 V AC
Output 1        : 53.28 V,  0.82 A   (12V/715W Max)
Power Usage     : 57.48 W
Airflow Direction : Front to Back

PowerSupply 2 information:
State           : Empty

System Power Usage : 57.48 W
Poll Interval     : 60 s
Change Threshold  : N/A
Change Action     : N/A

show fans
Slot-1 FanTray-1 information:
State:           Operational
NumFan:          2
Fan-1:           Operational at 4096 RPM
Fan-2:           Operational at 4488 RPM

Slot-1 FanTray-2 information:
State:           Operational
NumFan:          2
Fan-1:           Operational at 4028 RPM
Fan-2:           Operational at 4572 RPM

Slot-1 FanTray-3 information:
State:           Operational
NumFan:          2
Fan-1:           Operational at 4062 RPM
Fan-2:           Operational at 4509 RPM

show odometers

```

Field Replaceable Units	Service Days	First Recorded Start Date
-----	-----	-----
Switch : 5520-24W-EXOS	90	May-29-2020

History

This command was first available in ExtremeXOS 16.1.

Solid State Storage Device SSD-120 information was added in ExtremeXOS 30.3.

Secure Boot (trusted delivery) status information (5420, 5520, 5720 only) was added in beginning with ExtremeXOS 31.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

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

The output of this command displays the following information:

- TACACS+—The current state of TACACS+, enabled or disabled.
- TACACS+ Authorization—The current state of TACACS+ authorization, enabled or disabled.
- TACACS+ Accounting—The current state of TACACS+ accounting, enabled or disabled.
- TACACS+ Server Connect Timeout—The amount of time configured to detect and recover from a TACACS+ server failure.
- Primary TACACS+ Server—Describes information about the primary TACACS+ server, including:
 - The name of the primary TACACS+ server.
 - The IP address of the primary TACACS+ server.
 - The TCP port to use to contact the primary TACACS+ server.

- The IP address and VR used by the switch.
- The shared secret configured for the primary TACACS+ server.
- Secondary TACACS+ Server—Contains the same type of output as the primary TACACS+ server for the secondary TACACS+ server, if configured.
- TACACS+ Acct Server Connect Timeout—The amount of time configured to detect and recover from a TACACS+ accounting server failure.
- TACACS+ Accounting Server parameters, if configured. Contains the same type of output as the TACACS+ server for the TACACS+ accounting server(s), if configured.

Example

The following command displays TACACS+ client configuration and statistics:

```
show tacacs
```

The following is sample output from this command:

```
TACACS+: enabled
TACACS+ Authorization: enabled
TACACS+ Accounting : enabled
TACACS+ Server Connect Timeout sec: 3
Primary TACACS+ Server:
Server name   :
IP address    : 10.201.31.238
Server IP Port: 49
Client address: 10.201.31.65 (VR-Default)
Shared secret : qijxou
Secondary TACACS+ Server:
Server name   :
IP address    : 10.201.31.235
Server IP Port: 49
Client address: 10.201.31.65 (VR-Default)
Shared secret : qijxou
TACACS+ Acct Server Connect Timeout sec: 3
Primary TACACS+ Accounting Server:
Server name   :
IP address    : 10.201.31.238
Server IP Port: 49
Client address: 10.201.31.65 (VR-Default)
Shared secret : qijxou
Secondary TACACS+ Accounting Server:
Server name   :
IP address    : 10.201.31.235
Server IP Port: 49
Client address: 10.201.31.65 (VR-Default)
Shared secret : qijxou
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

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

The output of this command displays the following information:

- TACACS+ Accounting—The current state of TACACS+ accounting, enabled or disabled.
- TACACS+ Accounting Server Connect Timeout—The amount of time configured to detect and recover from a TACACS+ server failure.
- Primary TACACS+ Accounting Server—Describes information about the primary TACACS+ accounting server, including:
 - The name of the primary TACACS+ accounting server.
 - The IP address of the primary TACACS+ accounting server.
 - The TCP port to use to contact the primary TACACS+ accounting server.
 - The IP address and VR used by the switch.
 - The shared secret configured for the primary TACACS+ accounting server.
- Secondary TACACS+ Accounting Server—Contains the same type of output as the primary TACACS+ accounting server for the secondary TACACS+ accounting server, if configured.

Example

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

```
show tacacs-accounting
The following is sample output of this command:
TACACS+ Accounting : enabled
TACACS+ Acct Server Connect Timeout sec: 3
Primary TACACS+ Accounting Server:
Server name      :
IP address       : 10.201.31.238
Server IP Port   : 49
Client address   : 10.201.31.85 (VR-Default)
Shared secret    : qijxou
Secondary TACACS+ Accounting Server:Not configured
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show time

show time {**detail**}

Description

Shows the current time on the switch.

Syntax Description

detail	In addition to the current time, this option shows boot time, time zone, and the clock source.
---------------	--

Default

N/A.

Usage Guidelines

The date, time, time zone abbreviation, and offset from UTC appear in the following format:
<weekday_abbreviation> <month_abbreviation> <day_of_month> <HH:MM:SS> <timezone>(UTC<+|-><utc_offset>) <year>

If the **detail** option is selected, the following appears in the presented format:

```
Current Time: <current_time>
Boot Time: <boot_time>
Timezone: <DST_status> <GMT_offset> <timezone_name>
           <DST_in_effect>
           <DST_begin>
           <DST_end>
Source: [Local | NTP | SNTP] <ntp_peer_mode> <sntp_ntp_address>
```

<ntp_peer_mode> is only valid when the source is NTP. <ntp_peer_mode> can be any of the following:

- Symmetric Active
- Symmetric Passive
- Client, Server
- Broadcast Server
- Control, Private
- Broadcast Client

<sntp_ntp_address> is the peer address when source is NTP. <sntp_ntp_address> is the server address when source is SNTP.

The source is determined as follows:

If SNTP is enabled and active, SNTP is assumed to be the clock source. Otherwise, if NTP peer mode is not zero (unspecified) and the peer address is not 0.0.0.0, NTP is assumed to be the clock source. Otherwise, the clock source is local.

Example

The following example shows detailed time information on the switch:

```
# show time detail
Current Time:      Wed Oct 25 03:43:22 2017
Boot Time:        Sat Oct 21 04:12:53 2017
Timezone:         [Auto DST Disabled] GMT Offset: 0 minutes, name is UTC.
Source:           Local
```

History

This command was first available in ExtremeXOS 22.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show tech-support

```
show tech-support {all | area | {{bundle bundle_name} {monitor
num_iterations} {interval seconds}} {detail} {logto [file]}
```

Description

Displays the output of various show commands to assist in monitoring and troubleshooting the switch; use only in conjunction with Extreme Networks Technical support.

Syntax Description

all	Indicates all available show command output to be displayed.
<i>area</i>	Specifies one tech support area. For example, if you want to view <i>STP</i> information, enter <i>stp</i> .
detail	Specifies more detailed information.
logto [file]	Instructs the switch to log the show tech-support output into a file located in the switch's internal memory. The default file name is <i>show_tech.log.tgz</i> .
bundle	Selects collecting debug commands for specified functional areas.
<i>bundle_name</i>	Specifies the debug bundle name.

monitor	Selects executing a debug bundle multiple times.
<i>num_iterations</i>	Specifies the number of times to execute the debug bundle.
interval	Specifies setting the interval between successive executions of the debug bundle.
<i>seconds</i>	Sets the delay in seconds between successive intervals. The default is 300s.

Default

The default interval between multiple executions of a debug bundle is 300s.

Usage Guidelines



Note

Use this command only under the guidance of Extreme Networks Technical Support personnel to view your switch configurations and to troubleshoot the switch.

The `show tech-support` command displays the output of the following commands, among others:

- `ls /usr/local/tmp`
- `show bootprelay`
- `show configuration`
- `show dhcp-client state`
- `show diagnostics`
- `show management`
- `show memory`
- `show odometers`
- `show policy`
- `show port rxerror`
- `show port txerror`
- `show power`
- `show power budget`
- `show power controller`
- `show process`
- `show radius`
- `show session`
- `show switch`
- `show tacacs`
- `show version`
- `show vlan`

Information about the following areas is also displayed, among others:

- aaa
- bootp
- cli
- stp

If you enter the detail keyword, the following show output is displayed, among others:

- `show log`
- `show log configuration`
- `show log counters all`
- `show process detail`

This information can be useful for your technical support representative if you experience a problem.

Depending on the software version running on your switch, the configurations running on your switch, and the type of switch you have, additional or different show command and configuration output may be displayed.

Debug bundles are JSON files that provide predefined modules to common problems that can be executed when an issue occurs. You can search for a bundle that best matches your issue, and then execute that relevant bundle that matches your problem. For a list of bundles, use the command `show tech-support help`, or view information about this feature in the *Troubleshooting* chapter in the [Switch Engine 32.3 User Guide](#).

Example

The following example collects debugging information using a debug bundle named "l3" to run 5 times:

```
# show tech-support bundle l3 monitor 5
```

History

This command was first available in ExtremeXOS 10.1.

The command name and command syntax was modified in ExtremeXOS Release 15.4 from `show tech` to `show tech-support`.

Fabric attach mapping information was added in ExtremeXOS 22.4.

Debug bundle information was added in ExtremeXOS 30.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show tech-support collector

```
show tech-support collector { hostname | ip_address }
```


Description

This command displays tech-support configuration and report status.

Syntax Description

<i>hostname</i>	Specifies the host name of the collector.
<i>ip_address</i>	Specifies the IPv4 address of the collector.

Default

NA.

Usage Guidelines

This commands displays tech-support configuration and report status, such as last report stats, when the next cyclic report is sent if report frequency daily is configured, and when the next one time report is sent if one time report is scheduled with `run tech-support in hours` command. If you do not specify a hostname or IP address, this command displays configuration and report status for all existing collectors; otherwise it displays configuration and report status for the specified collector.

Example

The following is an example of the show output:

```
Tech Support Collector:      Enabled

Report Collector:           10.5.2.107
  TCP Port:                 9998
  Virtual Router:           VR-Mgmt
  Source IP Addr:           10.66.24.234
  SSL:                      Off
  Report Mode:              Automatic
  Report Data Set:           Summary
  Report Frequency:          Error-detected AND Daily at 0:00
  Last Report:              Successful
    Sent On:                Thu Feb 21 05:06:32 2013
    From:                   10.66.24.234:3546
    To:                     10.5.2.107:9998
    SSL:                    Off
    Error:                  None
  Next Cyclic Report:        Scheduled on Fri Feb 22 00:00:00 2013
  Next One Time Report:      Not Scheduled

Report Collector:           10.5.2.108
  TCP Port:                 800
  Virtual Router:           VR-Mgmt
  Source IP Addr:           10.66.24.234
  SSL:                      On
  Report Mode:              Automatic
  Report Data Set:           Detail
  Report Frequency:          Bootup AND Daily at 0:00
  Last Report:              Failed
    Sent On:                Thu Feb 21 05:06:32 2013
    From:                   10.66.24.234:3598
```

```
To: 10.5.2.107:800
SSL: On
Error: Socket time out
Next Cyclic Report: Scheduled on Fri Feb 22 00:00:00 2013
Next One Time Report: Not Scheduled
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show temperature

```
show temperature
```

Description

This command displays the current temperature of the power supply controllers and the switch.

On a stack, the command displays the current temperature of the switches in each slot.

Syntax Description

This command has no arguments or variables

Default

N/A.

Usage Guidelines

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

Use this command to display the temperature in Celsius and the current status of the standalone switch or all switches in a stack.

Bridge Port Extenders (BPEs)

Temperature information for bridge port extenders (BPEs) does not appear in this command. Use the `show vpex bpe {slot slot_num} {environment}` command instead.

Example

The following is sample output from a SummitStack of 8 nodes:

```
Slot-3 Stack.1 # show temperature
Field Replaceable Units          Temp (C)   Status   Min   Normal   Max
-----
Slot-1          :
Slot-2          : SummitX          34.50    Normal   -10    0-54    59
Slot-3          : SummitX          36.50    Normal   -10    0-66    67
Slot-4          :
Slot-5          :
Slot-6          :
Slot-7          :
Slot-8          :
Slot-3 Stack.2 #
```

History

This command was first available in an ExtremeXOS 10.1.

Information about the power controller(s), a component status column, and the minimum, normal, and maximum temperature ranges of the components was added to the output in ExtremeXOS 11.0.

Support for stacking was added to output in ExtremeXOS 12.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show tunnel

```
show [{tunnel} {tunnel_name} | tunnel {vr [vrname | all]} {payload-vr
payload_vrname} {detail}
```

Description

Displays system tunnel information for a specified tunnel or for all tunnels.

Syntax Description

<i>tunnel_name</i>	Specifies a tunnel name.
vr	Specifies showing tunnel information by VR.
<i>vrname</i>	Specifies showing tunnel information for a specific VR by name. The default is the context VR.
all	Specifies showing tunnels for all VRs.
payload-vr	Specifies showing tunnel information for payload VRs.
<i>payload_vrname</i>	Specifies showing tunnel information for a specific payload VR by name. If not specified, all tunnel payload VRs are shown.
detail	Shows detailed information for a tunnel.

Default

If not specified, the VR of current command context is used.

If not specified, all tunnel payload VRs are shown.

Usage Guidelines

The **tunnel** keyword is optional only when you specify a valid IPv6-in-IPv4 or IPv6-to-IPv4 tunnel name. The Total tunnels count in the display represents all tunnels on the switch.

Example

The following example displays system tunnel information for all tunnels:

```
# show tunnel
Name                               Type                               Flags
tunfour                           6in4 10.20.30.40 => 10.10.10.10  U
mytun                             GRE 1.1.1.2 => 1.1.1.1
Utunfive2                         6to4 10.20.30.40 => *.*.*.*      D
Total tunnels: 3
Flags: (U) Up / (D) Down / (a) Administratively Disabled
       (S) System Disabled (incompatible hardware)
```

The following example displays system tunnel information for tunnel "tunfour":

```
# show "tunfour"
Name                               Type                               Flags
tunfour                           6in4 10.20.30.40 => 10.10.10.10  U
Total tunnels: 2
Flags: (U) Up / (D) Down
```

The following example shows detailed information about tunnel "mytun1":

```
# show tunnel mytun1 detail
Tunnel Name : mytun
Tunnel Type : GRE
Flags : Administratively Disabled
Source Address : 1.1.1.1
Destination Address : 1.1.1.2
IPv4 Forwarding : Enabled
Virtual Router (VR) : vr-transport
Payload VR : vr-tenant
Interface Address : 2.0.0.1/24
```

The following example shows detailed information about tunnel "mytunnel" with TCP MSS adjustment configuration:

```
# show mytunnel detail
Tunnel Name      : mytunnel
Tunnel Type      : GRE
Flags            : Up
Source Address    : 1.1.1.1
Destination Address : 1.1.1.2
IPv4 Forwarding  : Enabled
Virtual Router (VR) : VR-Default
Payload VR       : VR-Default
Interface Address : 20.20.20.1/24
                  fe80::101:101/64
                  2001:100:100:100::1/64
```

```
TCP Adjust MSS      : IPv4: On      IPv6: On
TCP MSS Value       : IPv4: 9130    IPv6: 9110
```

The following example shows detailed information about tunnel "mytunnel" with an IP MTU value of 1476:

```
# show mytunnel detail
  Tunnel Name      : mytunnel
  Tunnel Type      : GRE
  Flags            : Up
  Source Address   : 1.1.1.1
  Destination Address : 1.1.1.2
  IPv4 Forwarding  : Enabled
  Virtual Router (VR) : VR-Default
  Payload VR       : VR-Default
  IP MTU           : 1476
  Interface Address : 20.20.20.1/24
                   fe80::101:101/64
                   2001:100:100:100::1/64
  TCP Adjust MSS   : IPv4: On      IPv6: On
  TCP MSS Value    : IPv4: 9130    IPv6: 9110
```

The following example shows detailed information about tunnel "mytunnel" with IPv6 transport in IPv4 GRE tunnel enabled:

```
# show mytunnel detail
  Tunnel Name      : mytunnel
  Tunnel Type      : GRE
  Flags            : Up
  Source Address   : 1.1.1.1
  Destination Address : 1.1.1.2
  IPv4 Forwarding  : Enabled
  IPv6 Forwarding  : Enabled
  Virtual Router (VR) : VR-Default
  Payload VR       : VR-Default
  Interface Address : 2.0.0.1/24
                   fe80::303:301/64
                   200e::1/64
                   200a::1/64
  IP MTU           : 1400
  TCP Adjust MSS   : IPv4: On      IPv6: On
  TCP MSS Value    : IPv4: 1410    IPv6: 1200
```

History

This command was first available in ExtremeXOS 11.2.

Options to specify VR/payload VR and detailed tunnel information were added in ExtremeXOS 31.2.

TCP MSS adjustment information was added in ExtremeXOS 31.6.

IPv6 transport in IPv4 GRE tunnel was added in ExtremeXOS 32.3.

Platform Availability

This command is available on the platforms listed for the IPv6 interworking feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

show twamp endpoint

```
show twamp endpoint {ipaddress ip {port udp_port}}
```

Description

This command displays the endpoint configured values, run-time data, and test session information.

Syntax Description

ipaddress	The endpoint IP address, either IPv4 or IPv6.
<i>ip</i>	Specifies the endpoint IP address.
port	Specifies the endpoint port.
<i>udp-port</i>	The UDP port the endpoint will listen on; range is 1025 – 65535.

Default

N/A.

Usage Guidelines

None.

Example

```
# show twamp endpoint ipaddress 19.1.1.100 port 5000

TWAMP Endpoint

Endpoint Information
  Local Address:          19.1.1.100  Listening Port:          5000
  Received Packets:      7948         Transmitted Packets:    7948
  Active Sessions:       5

Session created on Thu Nov 13 15:41:49 2014
  Peer Address:          19.1.1.2    Port:                    11001
  Sequence Number:      1575         Last recv'd packet:    84ms

Session created on Thu Nov 13 15:41:49 2014
  Peer Address:          19.1.1.2    Port:                    11002
  Sequence Number:      1555         Last recv'd packet:    16ms

Session created on Thu Nov 13 15:41:49 2014
```

History

This command was first available in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show twamp reflector

```
show twamp reflector
```

Description

This command displays the configured values and run-time information of the Session-Reflector and its endpoints.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

```
# show twamp reflector
TWAMP Session-Reflector Enabled

  Session Information
    Used: 165 of 200
    Timeout: 300 seconds
```

Endpoints	Port	Sessions	Rx Packets	Tx Packets
19.1.1.100	5000	5	3091	3091
19.1.1.100	5001	40	5521	5521
19.1.1.100	5002	40	4728	4728
19.1.1.100	5003	40	3916	3916
18.1.1.100	5000	40	9266	9266

```
Displayed 5 endpoints
```

History

This command was first available in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

Object Missing

This object is not available in the repository.

show upm event

show upm event *event-type*

Description

Displays UPM events of the specified type.

Syntax Description

<i>event-type</i>	Displays events of the specified type for all profiles. Valid values for event-type are:device-detectdevice-removeuser-authenticateuser-unauthenticated.
-------------------	--

Default

N/A.

Usage Guidelines

Use this command to display the following types of events:

- device-detect
- device-remove
- user-authenticate
- user-unauthenticated

Example

The following command displays device-detect events:

```
show upm event device-detect
```

The output of the command is similar to the following:

UPM Profile	PortList

profile1	3

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the Universal Port feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

show upm history

```
show upm history {profile profile-name | event upm-event | status [pass
| fail] | timer timer-name | detail}
```

Description

Displays (in a tabular column) a list of UPM profile events executed on the switch.

Syntax Description

<i>profile-name</i>	Displays UPM events for the specified profile.
<i>upm-event</i>	Displays UPM events that were triggered by the specified event.
status [pass fail]	Displays UPM events that meet the specified status, which is either pass or fail.
<i>timer-name</i>	Displays UPM events that were triggered by the specified timer.
detail	Displays additional detail on UPM events.

Default

N/A.

Usage Guidelines

This is useful for trouble shooting and testing.

Example

The following example shows what appears when no UPM events have been triggered:

```
# show upm history
-----
Exec   Event/           Profile           Port Status Time Launched
Id     Timer/ Log filter
-----
Number of UPM Events in Queue for execution: 0
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the Universal Port feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

show upm history exec-id

```
show upm history exec-id number
```

Description

Displays information about an instance of a UPM profile executed on the switch.

Syntax Description

<i>number</i>	Specifies the execution identifier for the event you want to view.
---------------	--

Default

N/A.

Usage Guidelines

To view the execution identifiers for which you can display information, enter the show upm history command.

Example

The following example shows information for the event identified as 8006:

```
# show upm history exec 8006
UPM Profile: p1
Event: User Request      , Time run: 2006-10-18 11:56:15
Execution Identifier: 8006      Execution Status: Pass
Execution Information:
1 # enable cli scripting
2 # set var EVENT.NAME USER-REQUEST
3 # set var EVENT.TIME 1161172575
4 # set var EVENT.PROFILE p1
5 # enable por 1:1
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the Universal Port feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

show upm profile

show upm profile *name*

Description

Displays a list of the UPM profiles on the system and some of their configuration information, or the contents of a specified profile.

Syntax Description

<i>name</i>	Displays the contents of the specified profile.
-------------	---

Default

N/A.

Usage Guidelines

To see a list of all UPM profiles on a switch, use the command without the name option. The resulting display shows the names of the profiles on the system and their status, active or disabled.

Use the name option to see the contents of a specific profile.

Example

The output of the command is similar to the following:

```
* BD-8808.36 # show upm profile
=====
UPM Profile Events Ports Flags
=====
p1 UPM Timer(t1) e
p1 device detect 1:1 e
p2 e
=====
Number of UPM Profiles: 2
Number of UPM Events in Queue for execution: 0
Flags: d - disabled, e - enabled
Event name: log-message(Log filter name) - Truncated to 20 chars
* BD-8808.37 # show upm profile "p1"
Created at : 2010-04-11 04:07:41
Last edited at : 2010-04-11 04:07:41
*****Profile Contents Begin*****
ena por 1:1
*****Profile Contents Ends*****
Profile State: Enabled
Profile Maximum Execution Time: 30
Events and ports configured on the profile:
=====
Event                               Port list/Log filter
=====
device-detect                        1:1
device-undetected                   :
```

```
user-authenticated      :  
user-unauthenticated    :  
=====
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the Universal Port feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

show upm timers

```
show upm timers
```

Description

Displays a list of the UPM timers on the system and some of their configuration information.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to see a list of all UPM timers on a switch. The resulting display shows timer configurations, the associated profile, and flags to indicate the timer status. Flags are defined below:

- a - active
- d - disabled
- p - periodic

Example

This command displays UPM timer configuration:

```
show upm timers
```

The output of this command is similar to the following:

```
* BD-8808.43 # show upm timers  
Current Time: 2006-10-16 14:03:55
```

```
-----
UPM Profile Flags Next Execution
Timer Name time
-----

t1 p1 ep 2006-10-16 14:04:00 (Every 10 secs)
timerA
-----

Flags: e - Profile is enabled, d: Profile is disabled
o -Timer is non-periodic, p - Timer is periodic
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the Universal Port feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

show var

```
show var {varname}
```



Note

This is a script command and operates only in scripts or on the command line when scripting is enabled with the following command: [enable cli scripting {permanent}](#).

Description

Displays the current session variables or the named variable.

Syntax Description

<i>varname</i>	Displays the variable specified, if present.
----------------	--

Default

N/A.

Usage Guidelines

Use this command to see the list of current session variables. The display includes the variable name and value.

Example

The output of this command is similar to the following:

```
Switch.7 # show var
-----
Count : 4
-----

-----
variableName          variableValue
-----
CLI.SESSION_TYPE      serial
CLI.USER              admin
STATUS                0
x                     66
-----
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show version

```
show version {detail | process name | images {partition partition}
```

Description

Displays the hardware serial and version numbers, the software version currently running on the switch, and (if applicable) the software version running on the power controllers.

Syntax Description

detail	Specifies display of platform name.
process	Specifies display of all of the processes on the switch.
<i>name</i>	Specifies display of a specific process on the switch.
images	Specifies the display of installed images.
<i>partition</i>	Specifies display of a specific partition (primary or secondary).

Default

N/A.

Usage Guidelines

The following describes the information displayed when you execute the `show version` or `show version detail` commands:

- **Part Number**—A collection of numbers and letters that make up the part number of the switch.
- **Serial Number**—A collection of numbers and letters that make up the serial number of the switch.



Note

For information about the physical location of the serial number on your switch, refer to the section that describes your specific switch model in the hardware documentation.

- **Image**—The ExtremeXOS software version currently running on the switch. If you have two software images downloaded on the switch, only the currently running ExtremeXOS version information is displayed. The information displayed includes the major version number, minor version number, a specific patch release, and the build number. The software build date is also displayed.
- **BootROM**—The BootROM version currently running on the switch. For ExtremeSwitching Universal switches, the BootROM version on both partitions appears.
- **FPGA**—The field-programmable gate array firmware version currently running on the switch.
- **PLD**—The programmable logic device firmware version currently running on the switch.
- **Certified Version**—Shows the OpenSSL FIPS and Linux versions built with the ExtremeXOS image.

Depending on the model of your switch and the software running on your switch, different version information may appear.

If you use the `process` option, you see the following information about the processes running on the switch:

- **Process Name**—The name of the process.
- **Version**—The version number of the process.
- **BuiltBy**—The name of the software build manager.
- **Link Date**—The date the executable was linked.

Example

The following command displays the hardware and software versions currently running on the switch:

```
# show version
Switch      : 800992-01-03 2011G-00108 Rev 03 BootROM: 2.2.1.3   IMG: 31.1.0.631

Image       : ExtremeXOS version 31.1.0.631 by release-manager
              on Tue Sep 1 13:58:24 EDT 2020
BootROM     : Default 2.2.1.3 Alternate 2.2.1.3
Diagnostics :
Certified Version : EXOS Linux 4.14.170, Extreme Networks FIPS Object Module 2.0.16a
Build Tools Version : exos-arm64-sdk-3.0.3.1.0
```

The following example shows SSD-120 device information:

```
# show version
Slot-1      : 800908-00-01 1847F-10076 Rev 01 IMG: 30.3.0.453
VIM5-4XE-1  : 800910-00-01 1845F-10106 Rev 01
XN-SSD-001-120-1: 800954-00-02 1917F-10030 Rev 02

Image       : ExtremeXOS version 30.3.0.453 by release-manager
              on Thu Jul 11 11:52:10 EDT 2019
```

```
Diagnostics :
Certified Version : EXOS Linux 4.14.123, Extreme Networks FIPS Object Module 2.0.16m
```

History

This command was first available in ExtremeXOS 10.1.

Solid State Storage Device SSD-120 information was added in ExtremeXOS 30.3.

BootROM version for ONIE switches (which do not have a BootROM) appears as "N/A" for compatibility with Chalet in ExtremeXOS 30.4.

BootROM version for both partitions was added in ExtremeXOS 31.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show virtual-network

```
show virtual-network {vn_name | vxlan vni vni | [vlan vlan_name | vman vman_name] }
```

Description

Shows the virtual networks that have been created.

Syntax Description

<i>vn_name</i>	Alphanumeric string identifying the virtual network.
vxlan	Displays VXLAN virtual networks.
vni	A unique 24-bit Virtual Network Identifier that is used in the VXLAN header to encapsulate tenant frames.
<i>vni</i>	Virtual Network Identifier value between 1 and 16,777,215.
vlan	Specifies <u>VLAN</u> .
<i>vlan_name</i>	Name of the tenant VLAN.
vman	Specifies VMAN.
<i>vman_name</i>	Name of the tenant VMAN.

Default

N/A.

Usage Guidelines

When no option is specified a summary output of all virtual networks appears. You can also query the switch using the virtual network name, VXLAN VNI or the tenant VLAN/VMAN name. With each of these options, detailed information about the virtual networks appears.

Example

The following example shows all virtual networks:

```
# show virtual-network
Virtual Network
=====
vn_blue          Encap      ID      Flags      Tenant VLAN
                  Encap      ID      Encap      Role (Replicator)
                  Flags
=====
vn_blue          VXLAN      101020   LRX         tenant-blue
                  Leaf (192.168.100.210)
vniname1         VXLAN      101021   LRX         tenant-1
                  Leaf (192.168.100.210)
vniname2         VXLAN      101022   LRX         tenant-2
                  Leaf (192.168.111.211)
vniwithveryveryveryverylongname  T-         tenant-3
                  VXLAN      101023   LRX         Leaf
=====
Flags: (T) OTM Configured, (V) OVSDB Configured
Encap Flags: (L) Local Endpoints Configured,
              (R) Remote Endpoints Associated,
              (X) Exclude Tag
-----

Total number of Virtual Networks : 3
Local Endpoints                  : 1.1.1.121 (VR-Default)
Network Ports [VXLAN]           : 1-128
Dynamic Virtual Networks         : Off
Multicast Group Range            : 232.1.1.0/255.255.255.0
[OR]
Multicast Group Range            : None
```

When multicast group is not configured, [OR] configured on a virtual network, [OR] automatically assigned:

```
# show virtual-network VNET1
Name                          : VNET1
Created By                     : CLI,
Tenant VLAN                    : v100
Local Endpoint                 : 4.4.4.4
Flooding                       : Multicast
Group IP                       : None
[OR]
Group IP                       : 232.1.1.1
[OR]
Group IP                       : 232.1.1.1 (Auto-assigned)
                               == VXLAN Information ==
VNI                           : 1 (0.0.1)
Group IP                      :
Inner Tag                     : Exclude
Remote Endpoints               : 1.1.1.1 (VR-Default)
                               : 2.2.2.2 (VR-Default)
                               : 3.3.3.3 (VR-Default)
                               : 232.1.1.1 (VR-Default)
                               == End VXLAN Information ==
```

History

This command was first available in ExtremeXOS 21.1.

VMAN option added in ExtremeXOS 22.1.

Dynamic virtual networks and dynamic VLAN information was added in ExtremeXOS 30.3.

Flag for dynamically created VRs (D) was added in ExtremeXOS 30.4.

Assisted replication information added in ExtremeXOS 31.1.

Multicast Group Range added in ExtremeXOS 31.2.

Platform Availability

This command is supported on ExtremeSwitching Universal switches and stacks with Universal switch slots.

show virtual-network remote-endpoint vxlan

```
show virtual-network {vn_name} remote-endpoint vxlan {vni vni}  
                {ipaddress ipaddress {vr vr_name}}
```

Description

Use this command to show the remote-endpoints that have been created.

Syntax Description

<i>vn_name</i>	Display remote endpoints only for the specified Virtual Network string.
vni	A unique 24-bit Virtual Network Identifier that is used in the VXLAN header to encapsulate tenant frames.
<i>vni</i>	Virtual Network Identifier value between 1 and 16777215.
ipaddress	IP address of the remote endpoint.
<i>ipaddress</i>	A remote endpoint IP address

Default

N/A.

Usage Guidelines

When no VXLAN name or VNI is specified a summary output of all remote endpoints appears. You need to supply the virtual network name or VXLAN VNI to query the remote endpoints associated to a specific virtual network. To display detailed information about a specific remote-endpoint, use the **ipaddress** option.

Example

```
# show virtual-network remote-endpoint vxlan
IP Address Virtual Router Next Hops Virtual Active/Total Networks
=====
200.1.1.1 VR-Default 2/2 2
200.1.1.2 VR-Default 2/2 1
200.1.1.3 VR-Default 2/2 1
200.1.1.4 VR-Default 2/2 1
200.200.200.200 VR-Default 0/0 0
=====
Total number of Remote Endpoints : 5

# show virtual-network vxlan
remote-endpoint ipaddress 3.3.3.3

IP Address      : 3.3.3.3                VRF : VR-Default
Learning       : Enabled                Stats Monitor : On
Health Check    : None

                == Next Hop Information ==
Gateway        : 22.0.0.1                MAC  : 00:04:96:97:f8:a6
Interface      : trunk2                  Port  : 1:41
                == End Next Hop Information ==

                == Virtual Network Associations ==
Virtual Network : vni1                    VNI : 1
Origin          : OSPF Learned,

Virtual Network : vni2                    VNI : 2
Origin          : OSPF Learned,
                == End Virtual Network Associations ==
# show virtual-network remote-endpoint vxlan ipaddress 100.1.1.1

IP Address      : 100.1.1.1                VRF : VR-Default
Admin State     : Enabled
Learning        : Enabled                Stats Monitor : Off

                == Next Hop Information ==
Gateway        : 50.50.50.1                MAC  : 00:04:96:9a:92:b7
Interface      : s1                      Port  : 1:45
                == End Next Hop Information ==

                == Virtual Network Associations ==
Virtual Network : vn1                      VNI : 1
Origin          : OSPF Learned,
                == End Virtual Network Associations ==
```

History

This command was first available in ExtremeXOS 21.1.

Information about the remote endpoint's administrative state was added in ExtremeXOS 22.4.

Platform Availability

This command is supported on ExtremeSwitching Universal switches and stacks with Universal switch slots.

show virtual-network statistics

```
show virtual-network { vn_name | remote-endpoint vxlan {ipaddress  
    ipaddress} {vr vr_name}} statistics {no-refresh}
```

Description

Use this command to display statistics (byte/packet counters) on a Virtual Network remote endpoint.

Syntax Description

<i>vn_name</i>	Display statistics only for the specified Virtual Network string.
ipaddress	IP address of the remote endpoint.
<i>ipaddress</i>	A remote endpoint IP address.
no-refresh	Page by page display without auto-refresh.

Default

VR-Default.

Usage Guidelines

N.A.

Example

To show statistics on an existing Virtual Network remote endpoint:

```
# show virtual-network remote-endpoint vxlan ipaddress  
10.10.10.146 statistics  
VXLAN Statistics Tue Oct 13 13:40:32 2015  
Virtual-Network Rx Total Rx Byte Tx Total Tx Byte  
Remote-endpoint Frames Count Frames Count  
=====
```

10.10.10.146	0	0	53496	0		
--------------	---	---	-------	---	--	--

```
=====
```

0->Clear Counters U->page up D->page down ESC->exit

To show statistics on all existing Virtual Network remote endpoints:

```
# show virtual-network remote-endpoint vxlan statistics  
VXLAN Statistics Tue Oct 13 13:42:51 2015  
Virtual-Network Rx Total Rx Byte Tx Total Tx Byte  
Remote-endpoint Frames Count Frames Count  
=====
```

10.10.10.146	0	0	53496	0		
10.10.10.148	7870	928660	0	0		
100.99.1.202	0	0	6917	816206		

```
=====
```

0->Clear Counters U->page up D->page down ESC->exit

History

This command was first available in ExtremeXOS 21.1.

Platform Availability

This command is supported on ExtremeSwitching Universal switches and stacks with Universal switch slots.

show virtual-router

```
show virtual-router {name}
```

Description

Displays information about VRs and VRFs.

Syntax Description

<i>name</i>	Specifies the name of a VR or VRF.
-------------	------------------------------------

Default

N/A.

Usage Guidelines

The output display differs for the following options:

- `show virtual-router`—displays information about all VRs and VRFs.
- `show virtual-router vr_name`—displays information about a user VR or [*VR-Default*](#).
- `show virtual-router vrf_name`—displays information about the named VRF.

Example

The following example displays the VR and VRF configurations on the switch:

```
# show virtual-router
-----
Virtual      Number of   Number of   Flags
Router      VLANs       Ports
-----
VR-Control   0           0           -----S46
VR-Default   32          18          boprimORS46
xvr          1           0           b-----F46
VR-Mgmt      1           0           -----S46
-----
Flags: Routing protocols configured on the Virtual Router
      (b) BGP, (i) ISIS, (m) MPLS, (o) OSPF, (O) OSPFv3, (p) PIM,
      (r) RIP, (R) RIPng.
Virtual Router type
      (D) Dynamically created user VR, (F) Non-VPN VRF, (N) VPN VRF,
      (S) System, (U) User.
```

```

Virtual Router admin state
  (-) Disabled (4) IPv4 Enabled, (6) IPv6 Enabled.

System Totals      :
Total VRs          :      4      Max VRs          : 1066
Total User VRs     :      0      Max User VRs      : 63
Total Non-VPN VRFs :      1      Max VRFs         : 1000
Total VPN VRFs     :      0      Total System VRs  : 3
Total Protocols    :      8      Max Protocols     : 64
Max IPv4 VLANs     :    4096    Max IPv6 VLANs    : 1024
Total IPv4 VLANs   :     13     Total IPv6 VLANs : 8
Active IPv4 VLANs  :     12     Active IPv6 VLANs : 8
Inactive IPv4 VLANs :      1     Inactive IPv6 VLANs : 0

```

The following example displays information about VR-Default:

```

# show vr VR-Default
Virtual Router      : VR-Default          Type : System
Description        : Default VR
Operational State   : Up
IPv4 Admin State    : Enabled             IPv6 Admin State : Enabled
IPv4 Route Sharing  : Disabled            IPv6 Route Sharing : Disabled
L3VPN SNMP Traps    : Disabled
Protocols Configured :
-----
Protocol   Process      Configuration   Protocol
Name       Name         Module Name     Instances
-----
BGP        bgp            bgp             1
OSPF       ospf           ospf            1
PIM        pim           pim             1
RIP        rip           rip             1
ISIS       isis          isis            1
OSPFv3     ospfv3        ospfv3          1
RIPng      ripng         ripng           1
-----
Port List      : 1-34
VLANs          : Default
Virtual Router Totals :
Total Protocols :      7      Max Protocols      :      8
Total Ports     :     34
Total VLANs     :      1
Total IPv4 VLANs :      0      Total IPv6 VLANs :      0
Active IPv4 VLANs :      0      Active IPv6 VLANs :      0
Inactive IPv4 VLANs :      0      Inactive IPv6 VLANs :      0

```

The following example displays detailed information for a VRF:

```

# show virtual-router xvr
Virtual Router      : xvr                  Type : Non-VPN VRF
IPv4 Admin State    : Enabled              IPv4 Forwarding   : Enabled
IPv6 Admin State    : Enabled              IPv6 Forwarding   : Enabled
Operational State   : Up
IPv4 Route Sharing  : Enabled              IPv6 Route Sharing : Disabled
Parent VR           : VR-Default
Protocols Configured :
-----
Protocol   Process      Configuration   Protocol
Name       Name         Module Name     Instances
-----
BGP        bgp            bgp-3           2
-----
VLANs      : xlan
Virtual Router Totals :

```

```

Total Protocols      :    1      Max Protocols      :    8
Total Ports          :    0
Total VLANs          :    1
Total IPv4 VLANs     :    1      Total IPv6 VLANs     :    0
Active IPv4 VLANs    :    1      Active IPv6 VLANs    :    0
Inactive IPv4 VLANs  :    0      Inactive IPv6 VLANs  :    0

```

The following example displays information for user VR region1:

```

# show virtual-router
-----
Virtual          Number of   Number of   Flags
Router           VLANs       Ports
-----
uservr-1         26          0      bo---m--U46
  xxx            1          0      b-----N46
VR-Control       0          0      -----S46
VR-Default       12         14      bopri-ORS46
VR-Mgmt          1          0      -----S46
-----

Flags : Virtual Router Type
      (S) System, (U) User, (F) Non-VPN VRF, (N) VPN VRF
      (L) Local-only for switch's local IP addresses
      : Virtual Router Admin State
      (-) Disabled (4) IPv4 Enabled, (6) IPv6 Enabled
      : Routing protocols configured on the virtual router
      (b) BGP, (i) ISIS, (m) MPLS, (o) OSPF, (p) PIM, (r) RIP,
      (O) OSPFv3, (R) RIPvng

System Totals :
Total VRs      :    5      Max VRs      : 1066
Total User VRs :    1      Max User VRs :   63
Total Non-VPN VRFs :    0      Max VRFs      : 1000
Total VPN VRFs :    1      Total System VRs :    3
Total Protocols :   10      Max Protocols :   64
Max IPv4 VLANs : 4096      Max IPv6 VLANs : 1024
Total IPv4 VLANs :   19      Total IPv6 VLANs :    9
Active IPv4 VLANs :   18      Active IPv6 VLANs :    9
Inactive IPv4 VLANs :    1      Inactive IPv6 VLANs :    0

Show virtual-router detail for a user created VR
# show virtual-router "uservr-1"
Virtual Router      : uservr-1                      Type : User
IPv4 Admin State    : Enabled      IPv4 Forwarding : Enabled
IPv6 Admin State    : Enabled      IPv6 Forwarding : Enabled
Operational State   : Up
IPv4 Route Sharing  : Enabled      IPv6 Route Sharing : Disabled
L3VPN SNMP Traps    : Disabled
Protocols Configured :
-----
Protocol   Process   Configuration   Protocol
Name       Name      Module Name     Instances
-----
BGP        bgp-5      bgp-5           2
OSPF        ospf-5     ospf-5          1
MPLS        mpls-5     mpls-5          1
-----

VRFs Configured :
-----
Virtual          Flags
Router
-----
xxx              b-----N46
-----

Flags : Virtual Router Type

```

(S) System, (U) User, (F) Non-VPN VRF, (N) VPN VRF
 (L) Local-only for switch's local IP addresses
 : Virtual Router Admin State
 (-) Disabled (4) IPv4 Enabled, (6) IPv6 Enabled
 : Routing protocols configured on the virtual router
 (b) BGP, (i) ISIS, (m) MPLS, (o) OSPF, (p) PIM, (r) RIP,
 (O) OSPFv3, (R) RIPvng

Route Exports into L3VPN (BGP):

VPN VRF Policy	Route Type	Flags	Priority
xxx None	Direct	EO	2048
vpn2 None	Static	EO	2048

Flags: (E) Export Enabled, (L) Export Operationally Off due to Low Memory,
 (O) Export Operationally On

VLANS : foo1, foo10, foo11,
 foo12,foo13, foo14,
 foo15,foo16, foo17,
 foo18,foo19, foo2,
 foo3, foo4, foo5,
 foo6, foo7, foo8,
 foo9, lan-uvr-1, lo,
 lo-uvr-1, v77, v88,
 vlan1, vlan2

Virtual Router Totals :

Total Non-VPN VRFs	: 0	Total VPN VRFs	: 1
Total Protocols	: 3	Max Protocols	: 8
Total Ports	: 0	Total VLANS	: 26
Total IPv4 VLANS	: 7	Total IPv6 VLANS	: 1
Active IPv4 VLANS	: 7	Active IPv6 VLANS	: 1
Inactive IPv4 VLANS	: 0	Inactive IPv6 VLANS	: 0

Show virtual router for a VPN VRF t16.2 # sh virtual-router "xxx"

Virtual Router : xxx Type : VPN VRF
 IPv4 Admin State : Enabled IPv4 Forwarding : Enabled
 IPv6 Admin State : Enabled
 IPv6 Forwarding : Enabled
 Operational State : Up
 IPv4 Route Sharing :
 Enabled IPv6 Route Sharing : Disabled
 Parent VR : uservr-1
 VPN ID :
 VPN RD : 1:1
 Export RT : 1:1
 Import RT : 1:1
 Protocols Configured :

Protocol Name	Process Name	Configuration Module Name	Protocol Instances
BGP	bgp-5	bgp-3	2

Route Exports into L3VPN (BGP):

VPN VRF Policy	Route Type	Flags	Priority
xxx	Direct	EO	2048


```

None
-----
Flags: (E) Export Enabled, (L) Export Operationally Off due to Low Memory,
      (O) Export Operationally On
VLANs      : xlan
Virtual Router Totals :
Total Protocols      :      1      Max Protocols      :      8
Total Ports          :      0      Total VLANs        :      1
Total IPv4 VLANs     :      1      Total IPv6 VLANs   :      0
Active IPv4 VLANs    :      1      Active IPv6 VLANs   :      0
Inactive IPv4 VLANs  :      0      Inactive IPv6 VLANs :      0

```

The current and configured values for **max-gateways** now apply to IPv6 gateway sets as well as IPv4, so these values are added to the output of `show ipconfig ipv6`.

```

# show virtual-router
Virtual Router          Number of      Number of Flags
                        VLANs          Ports
-----
VR-Boston              2              2      -----U46
VR-Control              0              0      -----S46
VR-Default              1             278      boprimORS46
VR-Mgmt                 1              0      -----S46
-----
Flags : Virtual Router Type
(S) System, (U) User, (F) Non-VPN VRF, (N) VPN VRF
(L) Local-only for switch's local IP addresses
: Virtual Router Admin State
(-) Disabled (4) IPv4 Enabled, (6) IPv6 Enabled
: Routing protocols configured on the virtual router
(b) BGP, (i) ISIS, (m) MPLS, (o) OSPF, (p) PIM, (r) RIP,
(O) OSPFv3, (R) RIPng
System Totals      :
Total VRs          :      4      Max      VRs          :    256
Total User VRs     :      1      Max      User VRs       :     63
Total Non-VPN VRFs :      0      Max      VRFs          :    190
Total VPN VRFs     :      0
Total System VRs   :      3
Total Protocols    :      8
Max Protocols      :     64
Max IPv4 VLANs     :    512      Max      IPv6 VLANs   :    512
Total IPv4 VLANs   :      0      Total IPv6 VLANs   :      0
Active IPv4 VLANs  :      0      Active IPv6 VLANs   :      0
Inactive IPv4 VLANs :      0      Inactive IPv6 VLANs :      0
Max Shared GWs (Cur) :    32      Max Shared GWs (Cfg) :    32

```

History

A command similar to this command was available in ExtremeXOS 10.1 (`show vr`).

This command was first available in ExtremeXOS 11.0.

Support for non-VPNVRFs was added in ExtremeXOS 12.5.

The show output for **max-gateways** was added in ExtremeXOS 15.3.

The "L" flag was added to signify local-only VRs in ExtremeXOS 22.6.

The "D" flag was added to show if a VR is created dynamically in ExtremeXOS 30.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show vlan

```
show vlan {virtual-router vr-name}
show [ { vlan } vlan_name | vlan vlan_list ] {ipv4 | ipv6}
show vlan [tag tag | detail] {ipv4 | ipv6}
show vlan ports
```

Description

Displays information about one or all VLANs.

Syntax Description

<i>vr-name</i>	Specifies a VR name for which to display summary information for all VLANs. If no VR name is specified, the software displays summary information for all VLANs in the current VR context. Note: User-created VRs are supported only on the platforms listed for this feature in the Switch Engine 32.3 Feature License Requirements document. On switches that do not support user-created VRs, all VLANs are created in VR-Default and cannot be moved.
<i>vlan_name</i>	Specifies a VLAN name for which to display detailed VLAN information.
<i>vlan_list</i>	Specifies a VLAN list of IDs to display detailed VLAN information.
<i>tag</i>	Specifies the 802.1Q tag of a VLAN for which to display detailed VLAN information.
detail	Specifies that detailed information should be displayed for all VLANs.
ipv4	Specifies IPv4.
ipv6	Specifies IPv6.
ports	Displays VLAN ports information.

Default

Summary information for all VLANs on the device.

Usage Guidelines



Note

To display IPv6 information, you must issue either the `show vlan detail` command or `show vlan` command with the name of the specified VLAN.

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 names, descriptions, 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.



Note

The switches display the Mgmt VLAN in [*VR-Mgmt*](#).

When an IPv6 address is configured for the VLAN, the system may display one of the following two address types in parentheses after the IPv6 address:

- Tentative
- Duplicate



Note

For information about IPv6 address types, see the [Switch Engine 32.3 User Guide](#).

You can display additional useful information on VLANs configured with IPv6 addresses by issuing the `show ipconfig ipv6 vlan vlan_name` command.

When a displayed VLAN is part of a PVLAN, the display includes the PVLAN name and type (which is network, non-isolated subscriber, or isolated subscriber).

When the displayed VLAN is configured for VLAN translation, the display provides translation VLAN information. If the displayed VLAN is a translation VLAN, a list of translation VLAN members appears. If the displayed VLAN is a member VLAN, the display indicates the translation VLAN to which the member VLAN belongs.

Ports that are dynamically added by MVRP are indicated by the symbol "H".

Example

The following example shows information about VLAN "vlan1":

```
# show vlan vlan1
VLAN Interface with name host created by user
  Admin State:           Enabled      Tagging:Untagged (Internal tag 4094)
  Description:           None
  Virtual router:        VR-Default
  IP Anycast:             Enabled
  IPv4 Forwarding:       Enabled
  IPv4 MC Forwarding:    Disabled
  Primary IP:            60.1.1.1/24
  Secondary IPs:         70.1.1.1/24
  Anycast IP:            80.1.1.1/24
```

```

IPv6 Forwarding:      Disabled
IPv6 MC Forwarding:   Disabled
IPv6:                 None
STPD:                 None
Protocol:             Match all unfiltered protocols
Loopback:             Disabled
NetLogin:             Disabled
QosProfile:           None configured
Egress Rate Limit Designated Port: None configured
Flood Rate Limit QosProfile: None configured
Suppress ARP:         Disabled
Suppress ND:          Disabled
Proxy ARP:            Entry required
Ports: 0.             (Number of active ports=0)

```

The following sample output shows general VLAN status:

```

# show vlan
Untagged ports auto-move: Off
-----
---
Name                VID  Protocol Addr      Flags
Proto  Ports  Virtual Active router      /Total
-----
---
Default             1  -----
ANY                                0/0
VR-Default  ext  4094 -----
ANY                                0 /12
VR-Default  Mgmt 4095 -----
ANY                                1/1
VR-Mgmt
-----
---
Flags : ((B) BFD Enabled, (c) 802.1ad customer VLAN, (C) EAPS Control VLAN,
        (d) Dynamically created VLAN, (D) VLAN Admin Disabled,
        (E) ESRP Enabled, (f) IP Forwarding Enabled,
        (F) Learning Disabled, (i) ISIS Enabled,
        (I) Inter-Switch Connection VLAN for MLAG, (k) PTP Configured,
        (l) MPLS Enabled, (L) Loopback Enabled, (m) IPmc Forwarding Enabled,
        (M) Translation Member VLAN or Subscriber VLAN, (n) IP Multinetting Enabled,
        (N) Network Login VLAN, (o) OSPF Enabled, (O) Virtual Network Overlay,
        (p) PIM Enabled, (P) EAPS protected VLAN, (r) RIP Enabled,
        (R) Sub-VLAN IP Range Configured, (s) Sub-VLAN, (S) Super-VLAN,
        (t) Translation VLAN or Network VLAN, (T) Member of STP Domain,
        (v) VRRP Enabled, (V) VPLS Enabled, (W) VPWS Enabled,
        (Y) Policy Enabled

Total number of VLAN(s) : 3

```

The following example displays VLAN ports information:

```

show vlan ports 1,2,3,4,5,6,7,8,9,10,11,12
-----
---
Name                VID  Protocol Addr      Flags
Proto  Ports  Virtual  Active router  /Total
-----
---
ext                4094
-----
ANY      0 /12  VR-Default
-----
---

```

```

Flags : ((B) BFD Enabled, (c) 802.1ad customer VLAN, (C) EAPS Control VLAN,
        (d) Dynamically created VLAN, (D) VLAN Admin Disabled,
        (E) ESRP Enabled, (f) IP Forwarding Enabled,
        (F) Learning Disabled, (i) ISIS Enabled,
        (I) Inter-Switch Connection VLAN for MLAG, (k) PTP Configured,
        (l) MPLS Enabled, (L) Loopback Enabled, (m) IPmc Forwarding Enabled,
        (M) Translation Member VLAN or Subscriber VLAN, (n) IP Multinetting Enabled,
        (N) Network Login VLAN, (o) OSPF Enabled, (O) Virtual Network Overlay,
        (p) PIM Enabled, (P) EAPS protected VLAN, (r) RIP Enabled,
        (R) Sub-VLAN IP Range Configured, (s) Sub-VLAN, (S) Super-VLAN,
        (t) Translation VLAN or Network VLAN, (T) Member of STP Domain,
        (v) VRRP Enabled, (V) VPLS Enabled, (W) VPWS Enabled,
        (Y) Policy Enabled

Total number of VLAN(s) : 3 (1 displayed)

show vlan ports 1 detail
  VLAN Interface with name ext created by user
  Admin State:      Enabled
  Tagging:Untagged (Internal tag 4094)
  Description:      None
  Virtual router: VR-Default
  IPv4 Forwarding: Disabled
  IPv4 MC Forwarding: Disabled
  IPv6 Forwarding: Disabled
  IPv6 MC Forwarding: Disabled
  IPv6:             None
  STPD:             None
  Protocol:         Match all unfiltered protocols
  Loopback:         Disabled
  NetLogin:         Disabled
  QosProfile:       None configured
  Flood Rate Limit QosProfile:      None configured
  Suppress ARP:     Disabled
  Proxy ARP:        Entry required
  Ports: 12.        (Number of active ports=0)
           Untag:    1,      2,      3,      4,      5,      6,      7,
                    8,      9,      10,     11,     12
Flags:  (*) Active, (!) Disabled, (g) Load Sharing port
        (b) Port blocked on the vlan, (m) Mac-Based port
        (i) Port inactivated on the vlan due to VXLAN configuration
        (a) Egress traffic allowed for NetLogin
        (u) Egress traffic unallowed for NetLogin
        (t) Translate VLAN tag for Private-VLAN
        (s) Private-VLAN System Port, (L) Loopback port
        (x) VMAN Tag Translated port
        (F) Dynamically added by Fabric Attach
        (G) Multi-switch LAG Group port
        (H) Dynamically added by MVRP
        (I) Dynamically added by IDM
        (N) Dynamically added by Netlogin
        (U) Dynamically added uplink port
        (V) Dynamically added by VM Tracking

```

The following example is the show output of a VLAN that was created dynamically by MVRP.

```

show vlan sys_vlan_0100
VLAN Interface with name sys_vlan_0100 created dynamically by MVRP
  Admin State:      Enabled      Tagging:      802.1Q Tag 100
  Description:      None
  Virtual router: VR-Default
  IPv4 Forwarding: Disabled
  IPv6 Forwarding: Disabled

```

```

IPv6:          None
STPD:          None
Protocol:      Match all unfiltered protocols
Loopback:      Disabled
NetLogin:      Disabled
QosProfile:    None configured
Egress Rate Limit Designated Port: None configured
Flood Rate Limit QosProfile:      None configured
Suppress ARP:  Disabled
Proxy ARP:     Entry required
Ports: 1.      (Number of active ports=1)
Tag: *2gH

Flags : ((B) BFD Enabled, (c) 802.1ad customer VLAN, (C) EAPS Control VLAN,
(d) Dynamically created VLAN, (D) VLAN Admin Disabled,
(E) ESRP Enabled, (f) IP Forwarding Enabled,
(F) Learning Disabled, (i) ISIS Enabled,
(I) Inter-Switch Connection VLAN for MLAG, (k) PTP Configured,
(l) MPLS Enabled, (L) Loopback Enabled, (m) IPmc Forwarding Enabled,
(M) Translation Member VLAN or Subscriber VLAN, (n) IP Multinetting Enabled,
(N) Network Login VLAN, (o) OSPF Enabled, (O) Virtual Network Overlay,
(p) PIM Enabled, (P) EAPS protected VLAN, (r) RIP Enabled,
(R) Sub-VLAN IP Range Configured, (s) Sub-VLAN, (S) Super-VLAN,
(t) Translation VLAN or Network VLAN, (T) Member of STP Domain,
(v) VRRP Enabled, (V) VPLS Enabled, (W) VPWS Enabled,
(Y) Policy Enabled

```

The following is the show output of a VLAN system with MAC-based VLAN port ingress filtering, indicated by the symbol "M":

```

# show vlan SYS_VLAN_0300
VLAN Interface with name SYS_VLAN_0300 created dynamically
Admin State:      Enabled      Tagging: 802.1Q Tag 300
Description:      None
Virtual router:   VR-Default
IP Anycast:       Disabled
IPv4 Forwarding:  Disabled
IPv4 MC Forwarding: Disabled
IPv6 Forwarding:  Disabled
IPv6 MC Forwarding: Disabled
IPv6:            None
STPD:            None
Protocol:        Match all unfiltered protocols
Loopback:        Disabled
NetLogin:        Disabled
QosProfile:      None configured
Flood Rate Limit QosProfile:      None configured
Suppress ARP:    Disabled
Suppress ND:     Disabled
Proxy ARP:       Entry required
Ports: 1.        (Number of active ports=0)
Untag: 1:2M
Flags: (*) Active, (!) Disabled, (g) Load Sharing port
        (b) Port blocked on the vlan, (m) Mac-Based port
        (M) Mac-Based port with ingress filtering on
        (i) Port inactivated on the vlan due to VXLAN configuration
        (a) Egress traffic allowed for NetLogin
        (u) Egress traffic unallowed for NetLogin
        (t) Translate VLAN tag for Private-VLAN
        (s) Private-VLAN System Port, (L) Loopback port
        (x) VMAN Tag Translated port
        (A) Dynamically added by Auto-peering
        (F) Dynamically added by Fabric Attach
        (G) Multi-switch LAG Group port

```

```
(H) Dynamically added by MVRP
(I) Dynamically added by IDM
(N) Dynamically added by Netlogin
(U) Dynamically added uplink port
(V) Dynamically added by VM Tracking
```

History

This command was first available in ExtremeXOS 10.1.

The IPv6 information was added in ExtremeXOS 11.2.

The netlogin information was added in ExtremeXOS 11.3.

The VR and administratively enabled/disabled information was added in ExtremeXOS 11.4.

The **tag** option was added in ExtremeXOS 12.4.4.

The OpenFlow status feature was added in ExtremeXOS 15.3.

The `vlan_list` variable was added in ExtremeXOS 16.1.

Showing the "untagged ports auto-move" status was added in ExtremeXOS 22.1.

Proxy ARP for reachable routes information added in ExtremeXOS 22.4.

Flag added for dynamically added by Fabric Attach information in ExtremeXOS 22.5.

Flag added for dynamically added by NetLogin information in ExtremeXOS 22.5.

Flag added for policy enabled in ExtremeXOS 22.5.

IP anycast status was added in ExtremeXOS 30.6.

IP anycast IP address information was added in ExtremeXOS 30.7.

Ingress filtering for MAC-based VLANs was added in ExtremeXOS 31.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show vlan description

```
show vlan description
```

Description

Displays a list of VLANs and VLAN descriptions.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following example displays the descriptions for all VLANs:

```
# show vlan description
-----
Name          VID  Description
-----
ctrl1         11   Control Vlan
ctrl2         102  Control Vlan 2
Default       1
v1            60   vlan 1
vplsVlan      3296 L2 VPN to home office
-----
Total number of VLAN(s) : 5
```

History

This command was first available in ExtremeXOS 12.4.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show vlan dhcp-config

```
show {vlan} vlan_name dhcp-config
```

Description

Displays the DHCP server's configuration for the specified VLAN.

Syntax Description

<i>vlan_name</i>	Specifies the VLAN of the DHCP server of interest.
------------------	--

Default

N/A.

Usage Guidelines

None.

Example

The following sample output displays the configuration of the DHCP server for the VLAN test:

```
#show vlan test dhcp-config
  DHCP Address Range           : 10.10.10.100->10.10.10.200
  Netlogin Lease Timer         : Not configured (Default =
10 seconds)
  DHCP Lease Timer             : Not configured (Default =
7200 seconds)
  Primary DNS Server           : 1.1.1.1
  Secondary DNS Server         : 2.2.2.2
  Ports DHCP Enabled           : 23
```

History

This command was first available in ExtremeXOS 11.0.

The output is modified to show primary and secondary DNS servers in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show vlan dhcp-address-allocation

```
show {vlan} vlan_name dhcp-address-allocation
```

Description

Displays the *DHCP* server's address allocation on a specified *VLAN*.

Syntax Description

<i>vlan_name</i>	Specifies the VLAN of the DHCP server of interest.
------------------	--

Default

N/A.

Usage Guidelines

None.

Example

The following command displays the configuration of the DHCP for the VLAN corp:

```
show vlan corp dhcp-address-allocation
```

The following is sample output from this command:

```
=====
IP                MAC                State      Lease Renewal Time
=====
10.0.0.2          00:02:03:04:05:00   Offered    0000:00:10
10.0.0.3          00:08:03:04:05:00   Assigned   0000:59:09
10.0.0.4          ee:1c:00:04:05:00   Assigned   0000:59:09
=====
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show vlan dynamic-vlan

```
show vlan dynamic-vlan
```

Description

Displays the configuration related to dynamically created VLANs.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays configuration related to dynamically created VLANs.

```
# sh vlan dynamic-vlan
Uplink Ports      : 12-15, 18-20
#
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show vlan eaps

```
show {vlan} vlan_name eaps
```

Description

Displays the *EAPS* configuration (control, partner, or not added to an EAPS domain) of a specific *VLAN*.

Syntax Description

<i>vlan_name</i>	Specifies a VLAN name.
------------------	------------------------

Default

N/A.

Usage Guidelines

Use this command to see if the specified VLAN is associated with an EAPS domain.

The output of this command displays whether the VLAN is a control or partner VLAN for an EAPS domain. This command also displays if the VLAN is not a member of any EAPS domain.

If a VLAN is a partner VLAN for more than one EAPS domain, all of the EAPS domains that the VLAN is a partner of appears in the output.

Example

The following sample output displays the EAPS configuration for the control VLAN orange in EAPS domain eaps1:

```
#show vlan orange eaps

Vlan is Control in following EAPS domain:
eaps1
```

The following sample output displays the EAPS configuration for the protected VLAN purple in EAPS domain eaps1:

```
#show vlan purple eaps

Vlan is Protected in following EAPS domain(s):eaps1
```

The following sample output displays information about the VLAN default not participating in EAPS:

```
#show vlan default eaps

Vlan has not been added to any EAPS domain
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show vlan fabric attach assignments

```
show [{vlan} vlan_name | vlan {vlan_id}] fabric attach assignments
```

Description

Displays VLAN to Network Service Identifier (NSI) information.

vlan	Specifies VLAN.
<i>vlan_name</i>	Specifies the name of the VLAN to show NSI information about.
<i>vlan_id</i>	Specifies the ID of the VLAN to show NSI information about.

Default

Not specifying a VLAN shows all VLAN to NSI assignment information.

Usage Guidelines

When an assignment is first received by the proxy, it is marked as "Pending". Mappings that are sent to the FA server are updated with the status returned by the server; either "Active" or "Rejected".

Example

The following example shows all VLAN to NSI assignment information:

```
# show vlan fabric attach assignments
Fabric Attach Mode: Server
VLAN  VLAN Name                               Type    ISID/NSI  Status
----  -
4012   NSI-4012Salem                               Dynamic  4012      Active
10     NSI-31247Salem                               Static   31247     Pending
1234   NSI-1234-SanJose                             Dynamic  1234      Rejected
1012   NSI-101010_RDU                               Static   101010    Active
```

History

This command was first available in ExtremeXOS 22.4.

The option **mapping** was changed to **assignments** in ExtremeXOS 22.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show vlan l2pt

```
show [vlan | vman] vlan_name {ports port_list} l2pt {detail}
```

Description

Displays the L2PT configuration and status of a service.

Syntax Description

vlan	Displays <u>VLAN</u> configuration.
vman	Displays VMAN configuration.
<i>vlan_name</i>	Specifies a VLAN name.
ports <i>port_list</i>	Displays the ports and port list separated by a comma (,) or dash (-).
detail	Displays the L2PT configuration and status in detail.

Default

Disabled.

Usage Guidelines

Use this command to display the L2PT configuration and status of a service.

Example

The following is an example of the `show vman ports l2pt` command:

```
# show vman cust2 ports 1:2,1:7 l2pt

Interface      L2PT Profile Name
-----
1:2            my_l2pt_prof
1:7            (none)
```

The following example illustrates the `show vman l2pt ports detail` command:

```
show l2vpn vpls cust1 l2pt

Interface      L2PT Profile Name
-----
123.112.111.111 my_other_l2pt_prof

# show vman cust2 ports 1:2,1:7 l2pt

Interface      L2PT Profile Name
-----
```

```

1:2          my_l2pt_prof
1:7          (none)

# show vman cust2 ports 1:2,1:7 l2pt detail

Port 1:2
  L2PT Profile Name      : my_l2pt_profile

  Protocol Filter Name   : filter1
    Destination Address: 01:80:C2:00:00:02
    Protocol Id Type    : etype
    Protocol Id Value   : 0x8902
    Field Offset        : 14
    Field Value         : 03
    Field Mask          : FF
    Action              : Tunnel
    CoS                 : Default
    DSCP                : 50
    DSCP Replace        : Yes
    Packets Transmitted: 2300
    Packets Received   : 2300

  Protocol Filter Name   : filter2
    Destination Address: 01:80:C2:00:00:00
    Protocol Id Type    : snap
    Protocol Id Value   : 0x4041
    Field Offset        :
    Field Value         :
    Field Mask          :
    Action              : Tunnel
    CoS                 : 7
    DSCP                : 50
    DSCP Replace        : Yes
    Packets Transmitted: 500
    Packets Received   : 500

Port 1:7
  L2PT Profile Name      : (none)

```

History

This command was first available in ExtremeXOS 15.5.

Support for DSCP was introduced in ExtremeXOS 31.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show vlan security

```
show [ {vlan} vlan_name | vlan vlan_list] security
```

Description

Displays the MAC limit-learning and lock-learning information for the specified VLAN.

Syntax Description

<code>vlan_name</code>	Specifies a VLAN name.
<code>vlan_list</code>	Specifies a VLAN list of IDs.

Default

N/A.

Usage Guidelines

The switch displays the following information:

- Port—Indicates the port on which the MAC address has been learned.
- Limit—Indicates that there is either a limited or unlimited amount of learned entries.
- State—Indicates that the current *FDB* entries for the port are permanent, no additional entries are learned, or that the port allows unlimited, dynamic learning.
- Learned—Specifies the number of learned entries.
- Blackholed—Specifies the number of blackholed entries.
- Locked—Specifies the number of locked entries.

Example

The following sample output displays the security setting of the *DHCP* server for the VLAN blue:

```
#show vlan blue security
Port      Limit    State    Learned    Blackholed    Locked
24        Unlimited Unlocked  0           0              0
```

History

This command was first available in ExtremeXOS 11.1.

The `vlan_list` variable was added in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show vlan statistics

```
show [{vlan} vlan_name |vlan vlan_list]statistics {no-refresh |refresh}
```

Description

Displays *VLAN* statistics at the VLAN level.

Syntax Description

<i>vlan_name</i>	Specifies to display VLAN statistics from the VLAN with this name.
<i>vlan_list</i>	Specifies a VLAN list of IDs.
no-refresh	Specifies that there is no continuous refresh. The prompt comes back to the user after fetching statistics once.

Default

N/A.

Usage Guidelines

This command displays statistics based on the sum of the statistics for individual ports. Use it to display the VLAN statistics monitored using the `configure ports [port_list|all] monitor vlan vlan_name | vlan_list {rx-only | tx-only}` command.

Example

The following command displays VLAN statistics:

```
# show vlan statistics
Vlan              Rx Total          Rx Byte          Tx
Total            Tx Byte
                  Frames
Count              Frames              Count
=====
Default          30251013
7326296          22034          901840
=====
```

If the VLAN contains ports that do not support a certain type of VLAN statistic, such as transmit statistics or byte counters, then a dash character (-) will be displayed in that column.

History

This command was first available in ExtremeXOS 12.0.

Support for ExtremeSwitching series switches was added in ExtremeXOS 12.5.

The **no-refresh** keyword was removed in ExtremeXOS 16.1.

The *vlan_list* variable was added in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show vlan stpd

```
show [ {vlan} vlan_name | vlan vlan_list] stpd {blocked-ports}
```


Description

Displays the *STP* configuration of the ports assigned to a specific *VLAN*.

Syntax Description

<i>vlan_name</i>	Specifies a VLAN name.
<i>vlan_list</i>	Specifies a VLAN list of IDs.

Default

N/A.

Usage Guidelines

If you have a VLAN that spans multiple *STPDs*, use this command to display the STP configuration of the ports assigned to that specific VLAN.

This command displays the following:

- STPD port configuration.
- STPD port mode of operation.
- STPD path cost.
- STPD priority.
- STPD state (root bridge, etc.).
- Port role (root designated, alternate etc.).
- STPD port state (forwarding, blocking, etc.).
- Configured port link type.
- Operational port link type.

If your VLAN has the same name as another component, for example an STPD, Extreme Networks recommends that you specify the identifying keyword as well as the name. If you do not specify the vlan keyword, the switch displays an error message similar to the following:

```
%% Ambiguous command: "show Test stpd"
```

In this example, to view the STPD state of VLAN Test, enter show vlan Test stpd.

If you enter a VLAN name that is not associated with an STPD or does not exist, the switch displays an error message similar to the following:

```
Failed to find vlan 'vlan1' or it has no STP domains configured on it
```

If this happens, check to make sure you typed the correct name of the VLAN and that the VLAN is associated with an STPD.

If your VLAN has a name unique only to that VLAN, the keyword vlan is optional.

Example

The following sample output displays the spanning tree configurations for the VLAN Default:

```
#show vlan default stpd
s0(enabled)  Tag: (none)  Ports: 8 Root/P/C: 80:00:00:01:30:94:79:00/-----/0
Port  Mode   State   Cost  Flags  Priority  Port ID Designated Bridge
1:1   802.1D LEARNING  19    eDbb-d- 16      8001    80:00:00:01:30:94:79:00
1:2   802.1D DISABLED  4      e----- 16      8002    00:00:00:00:00:00:00:00
1:3   802.1D DISABLED  4      e----- 16      8003    00:00:00:00:00:00:00:00
1:4   802.1D LEARNING  4      eDbb-d- 16      8004    80:00:00:01:30:94:79:00
1:5   802.1D LEARNING  4      eDbb-d- 16      8005    80:00:00:01:30:94:79:00
1:6   802.1D DISABLED  4      e----- 16      8006    00:00:00:00:00:00:00:00
1:7   802.1D DISABLED  4      e----- 16      8007    00:00:00:00:00:00:00:00
1:8   802.1D DISABLED  4      e----- 16      8008    00:00:00:00:00:00:00:00
----- Flags: -----
1:          e=Enable, d=Disable
2: (Port role)  R=Root, D=Designated, A=Alternate, B=Backup, M=Master, Y=Boundary
3: (Config type) b=broadcast, p=point-to-point, e=edge, a=auto
4: (Oper. type)  b=broadcast, p=point-to-point, e=edge
5:              p=proposing, a=agree
6: (partner mode) d=802.1d, w=802.1w, m=mstp
7:              i=edgeport inconsistency
8:              B = Boundary, I = Internal
```

History

This command was first available in ExtremeXOS 10.1.

The `vlan_list` variable was added in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show vm

```
show vm {vm_name | detail}
```

Description

Shows the operational status of Integrated Application Hosting (IAH) guest virtual machines (VMs).

Syntax Description

vm	Designates a virtual machine.
<i>vm_name</i>	Specifies the VM name to add show information about.
detail	Specifies showing detailed information.

Default

N/A.

Usage Guidelines

The IAH feature requires the Solid State Storage Device SSD-120.

Example

The following example shows information for VM "TPVM":

```
# show vm TPVM
VM Name: TPVM
  State:      Running
  Memory size: 4096 MB
  CPUs:       2
  Auto-start: Enabled
  VNC:        127.0.0.1:2 (Port 5902)
    Disk: vda
      Source: /mnt/vmstorage/TPVM/TPVM.img
      Disk bus type:      virtio
      Allocated size in bytes: 34359738368 (32.00 GB)
      Physical size in bytes: 4199288832 (3.91 GB)
      Read requests:      15947
      Bytes read:          349533184
      Write requests:     330
      Bytes written:       4293632
    Network interfaces:
      Attached switch ports: 33-34
    CPU utilization:
      User:                  10.00%
      System:                30.00%
    Memory utilization:
      Used:                   0.05 GB
      Available:              3.95 GB
```

The following example shows information for multiple VMs:

```
# show vm
S A Name                                |--Memory (GB)--| |---CPU---| Ports
                                Total  Used Avail  #   Sys% User%
- - - - -
* E TPVM                          4.00  0.05  3.95  2 10.00 30.00 33,34,Mgmt
  D tpvm2                        4.00   N/A   N/A  2   N/A   N/A   None

Status:      (*) Running
Auto-start:  (D) Disabled, (E) Enabled.
```

History

This command was first available in ExtremeXOS 30.3.

Clarifications were made to memory and CPU usage information in ExtremeXOS 30.5.

Disk bus type information was added in ExtremeXOS 30.7.

Platform Availability

This command is available on all platforms that support the Integrated Application Hosting (IAH) feature and have a Core or Premier license installed. For a list of platforms that support the IAH feature and for information about licenses, see the [Switch Engine 32.3 Feature License Requirements](#).

show vm guest interfaces

show **vm** *vm_name* **guest interfaces**

Description

Shows the interfaces of a running guest virtual machine (VM).

Syntax Description

vm	Designates a virtual machine.
<i>vm_name</i>	Specifies the VM name to show interfaces for.
guest	Specifies guest OS viewpoint.
interfaces	Specifies interfaces and addresses.

Default

N/A.

Usage Guidelines

The Integrated Application Hosting (IAH) feature requires the Solid State Storage Device SSD-120.

For this command to work, qemu guest agent must be running on the guest VM.

Example

The following example shows the interfaces for VM "vm1":

# show vm vm1 guest interfaces			
Name	MAC address	Protocol	Address

lo	00:00:00:00:00:00	ipv4	127.0.0.1/8
-	-	ipv6	::1/128
ens8	52:54:00:d4:88:59	ipv4	10.68.13.7/24
-	-	ipv6	fe80::5054:ff:fed4:8859/64
ens6	00:11:88:fe:ab:75	ipv4	192.168.1.5/24
-	-	ipv6	fe80::211:88ff:fefe:ab75/64
ens7	00:11:88:fe:ab:76	ipv4	192.168.0.5/24
-	-	ipv6	fe80::211:88ff:fefe:ab76/64

History

This command was first available in ExtremeXOS 30.3.

Platform Availability

This command is available on all platforms that support the Integrated Application Hosting (IAH) feature and have a Core or Premier license installed. For a list of platforms that support the IAH feature and for information about licenses, see the [Switch Engine 32.3 Feature License Requirements](#).

show vm virtual-interface

```
show vm vm_name virtual-interface {vf_name}
```

Description

Displays virtual interfaces of a VM (virtual machine).

Syntax Description

vm	Designates a virtual machine.
<i>vm_name</i>	Specifies the VM name to show information about.
virtual-interface	Specifies showing information about a virtual interface.
<i>vf_name</i>	Specifies virtual interface name. If you do not specify a name, all virtual interfaces are shown.

Default

If you do not specify a name, all virtual interfaces are shown.

Usage Guidelines

N/A.

Example

The following example shows virtual interfaces for VM "vm1":

```
# show vm vm1 virtual-interface
Insight  VLAN
VF Name   Port      Id    MAC Address
-----
vf-user11  33        11    52:54:00:d5:7c:2d
vf-user12  33        12    52:54:00:bb:47:a9
VF1-33    33        22    52:54:00:25:96:e2
VF1-34    34        N/A    52:54:00:2b:3c:e8
```

History

This command was first available in ExtremeXOS 30.5.

Platform Availability

This command is available on all platforms that support the Integrated Application Hosting (IAH) feature and have a Core or Premier license installed. For a list of platforms that support the IAH feature and for information about licenses, see the [Switch Engine 32.3 Feature License Requirements](#).

show vman

```
show vman {vman_name | vman_list} {ipv4 | ipv6}
show vman [tag tag | detail] {ipv4 | ipv6}
```

Description

Displays information about one or all VMANs.



Note

The information displayed for this command depends on the platform and configuration you are using.

Syntax Description

<i>vman_name</i>	Specifies that information is displayed for the specified VMAN.
<i>vman_list</i>	Specifies the vman list.
<i>tag</i>	Specifies a VMAN using the 802.1Q tag.
detail	Specifies that all information is displayed for each VMAN.
ipv4	Specifies IPv4.
ipv6	Specifies IPv6.

Default

Summary information for all VMANs on the switch.

Usage Guidelines

None.

Example

The following example displays a list of all the VMANs on the switch:

```
# show vman
-----
Name          VID  Protocol Addr      Flags                Proto  Ports  Virtual
Active router
/Total
-----
le1            4091 -----a          ANY    2 /2    VR-Default
le2            4090 -----a          ANY    0 /0    VR-Default
vml            4089 -----          ANY    0 /0    VR-Default
-----
Flags : (a) Learning Domain (C) EAPS Control vlan, (E) ESRP Enabled,
(f) IP Forwarding Enabled, (i) ISIS Enabled, (I) IP Forwarding lpm-routing Enabled,
(L) Loopback Enabled, (m) IPmc Forwarding Enabled,
(n) IP Multinetting Enabled, (N) Network LogIn vlan,
(o) OSPF Enabled, (p) PIM Enabled,
(P) EAPS protected vlan, (r) RIP Enabled, (T) Member of STP Domain,
```

```
(v) VRRP Enabled, (B) 802.1ah Backbone VMAN, (S) 802.1ah Service VMAN
Total number of vman(s) : 3
```

The following example displays information on a single VMAN named vman1:

```
# show vman vman1
VMAN Interface with name vman1 created by user
Admin State:      Enabled      Tagging:      802.1Q Tag 100
Virtual router: VR-Default
IPv4 Forwarding: Disabled
IPv6 Forwarding: Disabled
IPv6:             None
STPD:             None
Protocol:         Match all unfiltered protocols
Loopback:         Disabled
NetLogin:         Disabled
QosProfile:       None configured
Egress Rate Limit Designated Port: None configured
Flood Rate Limit QosProfile:      None configured
Ports:    2.      (Number of active ports=0)
Tag:       *1,      *2
CEP:       *3: CVID 20-29
           *4: CVID 10-19 translate 20-29
           *5: CVID 10-19 translate 20-29,CVID 30
Flags:     (*) Active, (!) Disabled, (g) Load Sharing port
(b) Port blocked on the vlan, (m) Mac-Based port
(a) Egress traffic allowed for NetLogin
(u) Egress traffic unallowed for NetLogin
(t) Translate VLAN tag for Private-VLAN
(s) Private-VLAN System Port, (L) Loopback port
(x) VMAN Tag Translated port
(G) Multi-switch LAG Group port
```

The Port CVID output was added in the display of `show vman vlan_name | detail` in ExtremeXOS 15.3.2:

```
VMAN Interface with name vml created by user
Admin State: Enabled      Tagging:      802.1Q Tag 1000
Description: None
Virtual router: VR-Default
IPv4 Forwarding: Disabled
IPv6 Forwarding: Disabled
IPv6:             None
STPD:             None
Protocol:         Match all unfiltered protocols
Loopback:         Disabled
NetLogin:         Disabled
QosProfile:       None configured
Egress Rate Limit Designated Port: None configured
Flood Rate Limit QosProfile:      None configured
Ports:    3.      (Number of active ports=3)
  Untag:      *21: Port CVID 5,
              *24: Port CVID 7,
  Tag:        *22
Flags:     (*) Active, (!) Disabled, (g) Load Sharing port
(b) Port blocked on the vlan, (m) Mac-Based port
(a) Egress traffic allowed for NetLogin
(u) Egress traffic unallowed for NetLogin
(t) Translate VLAN tag for Private-VLAN
(s) Private-VLAN System Port, (L) Loopback port
(x) VMAN Tag Translated port
(G) Multi-switch LAG Group port
```

The show vman detail command shows all the information shown in the `show vman vlan_name` command, but displays information for all configured VMANs.

History

- This command was first available in ExtremeXOS 11.0.
- Information on IEE 802.1ah was added in ExtremeXOS 11.4.
- The **tag** option was added in ExtremeXOS 12.4.4.
- Port CVID output was added in ExtremeXOS 15.3.2.
- The `vman_list` variable was added in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show vman eaps

```
show {vman} vman_name eaps
```

Description

Displays the EAPS domains to which the VMAN belongs.

Syntax Description

<code>vman_name</code>	Specifies the name of the VMAN for which EAPS information is to be displayed.
------------------------	---

Default

N/A.

Usage Guidelines

None.

Example

The following example displays a list of EAPS domains for the campus1 VMAN:

```
show vman campus1 eaps
```

History

This command was first available in ExtremeXOS 11.0.

Information on IEE 802.1ah was added in ExtremeXOS 11.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show vman ethertype

```
show vman ethertype
```

Description

Displays the ethertype information and secondary ethertype port_ list for VLANs, VMANs and PBBNs

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following example shows the command output on switches that support only VMANs:

```
vMan ethertype: 0x88a8
```

The following example shows the command output on switches that support PBBNs:

```
# show vman ethertype
vman ethertype : 0x88a8
bvlan ethertype: 0x88b5
```

The following example shows the command output when a secondary ethertype is configured with ports information:

```
# show vman ethertype
Vman Primary ethertype   : 0x9100
Vman Secondary ethertype : 0x8100
BVlan ethertype          : 0x88b5
Secondary ethertype ports : 6:2g 6:3
```

The letter g in the port list indicates that the port is a [LAG/Trunk](#) port, the details of which can be seen using the show port sharing command.

History

This command was first available in ExtremeXOS 11.0.

Information on IEE 802.1ah was added in ExtremeXOS 11.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show vm-tracking

show vm-tracking

Description

Displays the *XNV* feature configuration and the authenticated VM information.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays the XNV configuration and the authenticated VMs:

```
* Switch.51 # sh vm-tracking

-----
VM Tracking Global Configuration
-----
VM Tracking                               : Enabled
VM Tracking authentication order: nms vm-map local
VM Tracking nms reauth period   : 0 (Re-authentication disabled)
VM Tracking blackhole policy    : none
-----

Port                                     : 1:20
VM TRACKING                             : ENABLED

      Flags
MAC      APC      IP Address      Type      Value
-----
-----
```

```

-----
00:00:00:00:00:11 LBI      255.255.255.255    VM
                                   VPP      lvpp1
                                   IEP
                                   EEP
00:00:00:00:00:12 ---
                                   VM
                                   VPP
                                   IEP
                                   EEP
00:00:00:00:00:13 V---    30.30.30.30      VM      VMware-VM#2
                                   VPP      nvpp1
                                   IEP      al.pol
                                   EEP      a2.pol
-----

Flags :
(A)uthenticated : L - Local, N - NMS, V - VMMAp
(P)olicy Applied : B - All Ingress and Egress, E - All Egress, I - All Ingress
(C)ounter Installed: B - Both Ingress and Egress, E - Egress, I - Ingress

Type :
IEP - Ingress Error Policies
EEP - Egress Error Policies

Number of Network VMs Authenticated: 1
Number of Local VMs Authenticated  : 1
Number of VMs Authenticated        : 2

```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show vm-tracking local-vm

```
show vm-tracking local-vm {mac-address mac}
```

Description

Displays one or all of the VM entries in the local VM database.

Syntax Description

<i>mac</i>	Specifies the MAC address of a VM database entry that you want to display.
------------	--

Default

N/A.

Usage Guidelines

If you do not enter a MAC address with this command, the command displays all entries in the local VM database.

Example

The following command displays the local database VMs:

```
# show vm-tracking local-vm
```

MAC Address	IP Address	Type	Value
00:00:00:00:00:21		VM	
		VLAN Tag	100
		VR Name	VR-Default
		VPP	vpp1

Number of Local VMs: 1

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show vm-tracking network-vm

```
show vm-tracking network-vm
```

Description

Displays all of the VM entries in the network VM database.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays the configuration for all entries in the network VM database:

```
# show vm-tracking network-vm
```

MAC Address	IP Address	Type	Value
00:00:00:00:00:11	192.168.100.200	VM	KVM-VM-#101
		VPP	vpp300
00:01:02:03:04:06	192.168.100.201	VM	VM #200
		VPP	vpp201

```
Number of Network VMs: 2
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show vm-tracking nms

```
show vm-tracking nms server {primary | secondary}
```

Description

Displays the *RADIUS* client configuration and operating statistics for one or both NMS servers.

Syntax Description

primary secondary	Specifies whether you are displaying the primary or secondary NMS server information.
-----------------------------------	---

Default

If you do not specify primary or secondary, the default action is to display both the primary and secondary NMS server configurations.

Usage Guidelines

None.

Example

The following command displays the RADIUS client information for the primary and secondary NMS servers:

```
# show vm-tracking nms server
VM Tracking NMS (RADIUS): enabled
```

```
VM Tracking Radius server connect time out: 3 seconds
Primary VM Tracking NMS server:
Server name      :
IP address       : 10.127.5.221
Server IP Port   : 1812
Client address   : 10.127.10.173 (VR-Mgmt)
Shared secret    : pmckmtpq
Access Requests  : 0                Access Accepts      : 0
Access Rejects   : 0                Access Challenges   : 0
Access Retransmits: 0                Client timeouts     : 0
Bad authenticators: 0                Unknown types       : 0
Round Trip Time  : 0
Secondary VM Tracking NMS server:
Server name      :
IP address       : 10.127.5.223
Server IP Port   : 1812
Client address   : 10.127.10.173 (VR-Mgmt)
Shared secret    : rjgueogu
Access Requests  : 0                Access Accepts      : 0
Access Rejects   : 0                Access Challenges   : 0
Access Retransmits: 0                Client timeouts     : 0
Bad authenticators: 0                Unknown types       : 0
Round Trip Time  : 0
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show vm-tracking port

```
show vm-tracking port port_list
```

Description

Displays the *XNV* feature configuration for the specified port and information for all VMs authenticated on the port.

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
------------------	---

Default

N/A.

Usage Guidelines

None.

Example

The following command displays the XNV configuration for port 1:20 and the authenticated VMs:

```
# sh vm-tracking port 1:20
-----
      VM Tracking Global Configuration
-----
VM Tracking                               : Enabled
VM Tracking authentication order          : nms vm-map local
VM Tracking nms reauth period             : 0 (Re-authentication disabled)
VM Tracking blackhole policy              : none
-----

Port                                     : 1:20
VM Tracking                             : Enabled
VM Tracking Dynamic VLAN                  : Enabled

      Flags
-----
MAC      AP      IP Address      Type      Value
-----
00:00:00:00:00:11 LBI      255.255.255.255 VM
                                           VLAN Tag      100
                                           VR Name      VR-Default
                                           VPP          lvpp1
                                           IEP
                                           EEP
-----

Flags :
  (A)uthenticated : L - Local, N - NMS, V - VM MAP
  (P)olicy Applied : B - All Ingress and Egress, E - All Egress, I - All Ingress
  (C)ounter Installed : B - Both Ingress and Egress, E - Egress, I - Ingress
  All Ingress and Egress, E - All Egress, I - All Ingress
Type :
  IEP - Ingress Error Policies      EEP - Egress Error Policies

Number of Network VMs Authenticated: 0
Number of Local VMs Authenticated  : 1
Number of VMs Authenticated        : 1
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show vm-tracking repository

```
show vm-tracking repository {primary | secondary}
```

Description

Displays the FTP file synchronization configuration for NVPP and VM MAP files.

Syntax Description

primary secondary	Specifies whether you are displaying the primary or secondary FTP server configuration.
-----------------------------------	---

Default

If you do not specify primary or secondary, the default action is to display both the primary and secondary FTP server configurations.

Usage Guidelines

None.

Example

The following command displays the configuration for the primary and secondary FTP servers:

```
# show vm-tracking repository
Primary VM-Map FTP server:
Server name:
IP address      : 10.100.1.200
VR Name         : VR-Mgmt
Refresh-interval: 600 seconds
Path Name       : /pub (default)
User Name       : anonymous (default)
Secondary vm-map FTP server: Unconfigured
Last sync       : 16:35:15      Last sync server : Primary
Last sync status: Successful
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show vm-tracking vpp

```
show vm-tracking vpp {vpp_name}
```

Description

Displays the configuration of one or all VPPs.

Syntax Description

<i>vpp_name</i>	Specifies the name of an existing local VPP.
-----------------	--

Default

All.

Usage Guidelines

You can only specify local VPPs with this command. If you do not enter a VPP name with this command, the command displays all local and network VPPs.

Example

The following command displays the configuration of all VPPs:

# show vm-tracking vpp		
VPP Name	Type	Value

nvpp1	origin	network
	counters	ingress-only
	VLAN Tag	200
	VR Name	VR-Default
	ingress	ingLocal1.pol(1)
		ingLocal2.pol(2)
	egress	egrLocal1.pol(1)
		egrLocal2.pol(2)
lvpp1	origin	local
	counters	egress-only
	VLAN Tag	100
	VR Name	VR-Default
	ingress	ing1.pol(1)
	egress	egr1.pol(1)
		egr2.pol(2)
Number of Local VPPs : 1		
Number of Network VPPs: 1		

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show vpex

show vpex

Description

Shows the status of Extended Edge Switching VPEX mode, auto-configuration (partial automation) and Zero Touch Provisioning (ZTP) (full automation) status, ring re-balancing status, and if applicable, switch port to bridge port extender (BPE) slot assignments.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

For ZTP (full automation) status, the following states can appear:

- **Initializing**—Initializing.
- **Searching for BPEs**—Extended Edge Switching-capable switch (CB) is detected and system is waiting for a BPE to be attached in order to continue, or for the configuration to be saved in order to finish. The CB can remain in this state indefinitely.
- **Searching for MLAG Peer**—Extended Edge Switching-capable switch (CB) is detected and system is looking for an peer. The CB can remain in this state for up to four minutes.
- **Nothing Provisioned – Existing configuration present**—Full automation is not running because it has detected an existing configuration.
- **Nothing Provisioned – Disabled**—Full automation is disabled by running the `terminate vpex ztp` command or by ExtremeCloud IQ - Site Engine.
- **Provisioning Complete**—Full automation has finished creating the Extended Edge Switching configuration.

Example

The following example shows VPEX mode, auto-configuration, automatic upgrading, and ring rebalancing are enabled, and that switch port 1:20 is attached to a BPE assigned to slot number 100:

```
# show vpex
Virtual Port Extender:  Enabled
Auto-Configuration:    Enabled
Auto-Upgrade:          Enabled
Ring rebalancing:      Auto
Zero Touch Provisioning: Nothing Provisioned - Existing configuration present

Cascade
Port/MLAG Id    Slot
=====
1:20            100
```

The following example shows that VPEX mode is disabled:

```
# show vpex
Virtual Port Extender:  Disabled
```

History

This command was first available in ExtremeXOS 22.5.

Auto-configuration status information was added in ExtremeXOS 22.6.

Ring re-balancing and Zero Touch Provisioning (ZTP) information was added in ExtremeXOS 22.7.

Automatic upgrading status was added in ExtremeXOS 30.5.

Platform Availability

This command is available on ExtremeSwitching 5420 and 5520 series switches.

show vpex auto-configuration

```
show vpex auto-configuration
```

Description

Shows the status of Extended Edge Switching auto-configuration (partial automation) and, if applicable, switch port to bridge port extender (BPE) slot assignments.

Syntax Description

vpex	Specifies Virtual Port Extender (VPEX).
auto-configuration	Specifies the VPEX Auto-Configuration of cascade ports and VPEX slots.

Default

N/A.

Usage Guidelines

None.

Example

The following example displays output when there is no MLAG peer configured:

```
# show vpex auto-configuration
VPEX Auto-Configuration: Enabled without MLAG support
Virtual MLAG ID Configuration: Cascade
```

The following example displays output when an MLAG peer is configured:

```
# show vpex auto-configuration
VPEX Auto-Configuration: Enabled with MLAG support
Virtual MLAG ID Configuration: Cascade
```

The following example displays output when an MLAG peer is configured and an MLAG peer is down:

```
# sh vpex auto-configuration
VPEX Auto-Configuration: Disabled
Virtual MLAG ID Configuration: Cascade
```

History

This command was first available in ExtremeXOS 31.7

Platform Availability

This command is available on ExtremeSwitching 5420 and 5520 series switches.

show vpex bpe

```
show vpex bpe
```

Description

Shows attached bridge port extender's (BPE's) cascade port, slot assignment number, and MAC addresses.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Example

The following example shows two BPEs attached to ports 111 and 102:51 assigned to slots 102 and 103, respectively:

Slot	Cascade		Model	PE		Description
	Port/MLAG	Id		MAC	Address	
102	111		V400-48t-10GE4	d8:84:66:f2:a4:3f	none	
103	102:51		V400-48t-10GE4	d8:84:66:f2:a8:98	none	

History

This command was first available in ExtremeXOS 22.5.

MLAG ID information was added in ExtremeXOS 22.7.

Platform Availability

This command is available on ExtremeSwitching 5420 and 5520 series switches.

show vpex bpe cpu-utilization

```
show vpex bpe {slot slot_num} cpu-utilization
```

Description

Displays the CPU utilization on the specified bridge port extenders BPE(s).

Syntax Description

vpex	Specifies VPEx.
bpe	Specifies BPE.
slot	Specifies a specific BPE by slot number.
<i>slot_num</i>	Specifies a specific BPE by slot number.
cpu-utilization	Specifies showing CPU information.

Default

N/A.

Example

The following example shows CPU information for all BPEs attached to the control bridge (CB):

```
# show vpex bpe cpu-utilization
Slot: 100 Cascade Port: 1:44 MAC: d8:84:66:88:98:06
CPU Utilization:
  5 seconds: 28 %
  1 minute:  16 %
  5 minutes: 16 %

Slot: 101 Cascade Port: 1:45 MAC: d8:84:66:88:98:07
CPU Utilization:
  5 seconds: 29 %
  1 minute:  17 %
  5 minutes: 18 %
```

History

This command was first available in ExtremeXOS 22.5.

Platform Availability

This command is available on ExtremeSwitching 5420 and 5520 series switches.

show vpex bpe environment

```
show vpex bpe {slot slot_num} {environment}
```

Description

Shows power supply, temperature, and fan status for bridge port extenders (BPE).

Syntax Description

vpex	Specifies Virtual Port Extender (VPEX).
bpe	Bridge port extender (BPE).
slot	Specifies a specific port extender by slot number.
<i>slot_num</i>	Specifies a specific port extender by slot number.
environment	Specifies port extender environment information, including PSU, fan, and temperature information.

Default

N/A.

Example

The following example shows environment information for all BPEs attached to the switch. In this case, there are two BPEs: one attached on port 1:45 and assigned to slot 100, and another attached on port 1:45 and assigned to slot 101:

```
# show vpex bpe environment

BPE Environment Information

Cascade Port: 1:41   Slot: 100   MAC: d8:84:66:88:98:74
  Description: Floor1
  Serial #:      800879-00-00 1716A042010
  Main PSU:      Powered On
  Backup PSU:     Empty
  Fan-1:         Operational at 5434 RPM
  Sensor-1:      Operational      Temp: Normal
  Sensor-2:      Operational      Temp: Normal
  Sensor-3:      Operational      Temp: Normal
  Sensor-4:      Operational      Temp: Normal
Cascade Port: 1:45   Slot: 101   MAC: d8:84:66:88:97:fa
  Description: Floor2
  Serial #:      00.00.01 1705D-10009
  Main PSU:      Powered On
  Backup PSU:     Empty
  Fan-1:         Operational at 10621 RPM
  Sensor-1:      Operational      Temp: Normal
  Sensor-2:      Operational      Temp: Normal
  Sensor-3:      Operational      Temp: Normal
  Sensor-4:      Operational      Temp: Normal
```

History

This command was first available in ExtremeXOS 22.5.

Platform Availability

This command is available on ExtremeSwitching 5420 and 5520 series switches.

show vpex bpe statistics

```
show vpex bpe {slot slot_num} {statistics} {detail}
```

Description

Shows bridge port extender (BPE) statistics.

Syntax Description

vpex	Specifies Virtual Port Extender (VPEX).
bpe	Bridge port extender (BPE).
slot	Specifies a specific BPE by slot number.
<i>slot_num</i>	Specifies a specific BPE by slot number.
statistics	Specifies BPE statistics information and control plane protocol.
detail	Shows detailed BPE statistics information and control plane protocol. When detail is excluded, only non-zero statistics appear; when detail is included, non-zero and zero statistics appear.

Default

N/A.

Usage Guidelines

This command displays information on a per-BPE basis (cascade ports can be LAG members). If the **detail** option is not used, only non-zero counts appear.

Example

The following example shows non-detailed statistics information for all BPEs. In this case, there is only one BPE attached at port: 1:45 and assigned to slot 100:

```
# show vpex bpe statistics
BPE Statistics

BPE: d8:84:66:88:98:2b   Cascade Port: 1:45   Slot: 100
Description: Floor1

  txRequests      : 23432      rxResponses      : 23432
  txResponses     : 52         rxRequests       : 52
  rxrqErrors      : 0          rxrspErrors      : 0

Standard Statistics
-----
                                TX Requests                RX Responses
cspOpen           : 1                                     1
echanRegister     : 2                                     2
orgSpecific       : 23429                                23429
                                TX Responses                RX Requests
```

```

cspOpen      : 1          1
extPortCreate : 47        47
statusParamSet : 4        4

Organizational Specific Statistics
-----
TX Requests                                RX Responses
frmSizeLimSet : 47          47
getBulk       : 15488       15488
qCountersGet  : 7789        7789
qTrafShaperSet : 51         51
defPrtQoSParSet : 47        47
dtaCtlTrfPriSet : 6         6
ipInterfaceSet : 1          1

RX Requests                                TX Responses
No Valid Statistics

```

The following example shows detailed statistics information for all BPEs. In this case, there is only one BPE attached at port: 1:45 and assigned to slot 100:

```

# show vpex bpe statistics detail
BPE Statistics

BPE: d8:84:66:88:98:2b   Cascade Port: 1:45   Slot: 100

txRequests      : 23630      rxResponses      : 23630
txResponses     : 52         rxRequests       : 52
rxrqErrors      : 0          rxrspErrors      : 0

Standard Statistics
-----
TX Requests                                RX Responses
cspOpen         : 1          1
extPortCreate   : 0          0
extPortDelete   : 0          0
portParamSet    : 0          0
portParamGet    : 0          0
statusParamSet  : 0          0
echanRegister   : 2          2
echanRegGet     : 0          0
statisticsGet   : 0          0
transitDelaySet : 0          0
objectGet       : 0          0
objectSet       : 0          0
cnParamGet      : 0          0
cnParamSet      : 0          0
orgSpecific     : 23627      23627
unsupported      : 0          0

TX Responses                                RX Requests
cspOpen         : 1          1
extPortCreate   : 47         47
extPortDelete   : 0          0
portParamSet    : 0          0
portParamGet    : 0          0
statusParamSet  : 4          4
echanRegister   : 0          0
echanRegGet     : 0          0
statisticsGet   : 0          0
transitDelaySet : 0          0
objectGet       : 0          0
objectSet       : 0          0

```



```

cnParamGet      : 0          0
cnParamSet      : 0          0
orgSpecific     : 0          0
unsupported      : 0          0

Organizational Specific Statistics
-----
                                TX Requests                                RX Responses
reload           : 0          0
swapImg          : 0          0
imageInfoGet     : 0          0
trstdPrtInfGet   : 0          0
poeParamSet      : 0          0
portCfgSet       : 0          0
portOpStatGet    : 0          0
frmSzeLimSet     : 47         47
frmSzeLimGet     : 0          0
lagBalAlgSet     : 0          0
ecpParamSet      : 0          0
cbMsgRefl        : 0          0
envGet           : 0          0
closeNotif       : 0          0
getBulk          : 15620      15620
localMirrSet     : 0          0
debugInfoGet     : 0          0
qCountersGet     : 7857       7857
qTrafShaperSet   : 51         51
cpuUtilGet       : 0          0
diagSet          : 0          0
addtlCapSet      : 0          0
defPrtQoSParSet  : 47         47
dtaCtlTrfPriSet  : 6          6
ipInterfaceSet   : 1          1
lagSet           : 0          0
portCntrClear    : 0          0
poeFltReasGet    : 0          0
eeeSet           : 0          0
poeGlobParmSet   : 0          0
tftpDownload     : 0          0
sfpNotifSet      : 0          0
diagInfoGet      : 0          0
unsupported       : 0          0

                                RX Requests                                TX Responses
POE Notif        : 0          0
cbMsgRefl        : 0          0
sfpNotif         : 0          0
unsupported       : 0          0

```

History

This command was first available in ExtremeXOS 22.5.

Platform Availability

This command is available on ExtremeSwitching 5420 and 5520 series switches.

show vpex bpe version detail

```
show vpex bpe { slot slot_num } version detail
```

Description

Displays detailed bridge port extender (BPE) image information for the specified BPE(s).

Syntax Description

vpex	Specifies Virtual Port Extenders (VPEX).
bpe	Specifies BPE.
slot	Specifies a specific BPE by slot number.
<i>slot_num</i>	Specifies a specific BPE by slot number.
version detail	Specifies detailed image version information for the specified BPEs.

Default

N/A.

Example

The following example show detailed image version information for all attached BPEs:

```
# show vpex bpe version detail

BPE Version Information

Cascade Port: 1:47 Slot: 100 MAC: d8:84:66:88:98:02
Active-image:
  Version:      1.0.0.35
  Date:         12-Jun-2017
  Time:         11:15:03
  MD5 Digest:   2bec22a139c0e0d243d9eed192ce54e4
Inactive-image:
  Version:      1.0.0.32
  Date:         24-May-2017
  Time:         19:07:31
  MD5 Digest:   04a0e44f58a9e0810ea16fc034bea462

Cascade Port: 1:45 Slot: 101 MAC: d8:84:66:88:98:05
Active-image:
  Version:      1.0.0.35
  Date:         12-Jun-2017
  Time:         11:15:03
  MD5 Digest:   2bec22a139c0e0d243d9eed192ce54e4
Inactive-image:
  Version:      1.0.0.32
  Date:         24-May-2017
  Time:         19:07:31
  MD5 Digest:   04a0e44f58a9e0810ea16fc034bea462
```

History

This command was first available in ExtremeXOS 22.5.

Platform Availability

This command is available on ExtremeSwitching 5420 and 5520 series switches.

show vpex ports

```
show vpex ports {ports_list}
```

Description

Shows information about the bridge port extender (BPE) attached to the specified cascade ports.

Syntax Description

vpex	Specifies Virtual Port Extender (VPEX).
ports	Specifies the cascade port attached to the BPE.
<i>port_list</i>	Specifies the cascade port attached to the BPE.

Default

If you do not specify any ports, information for all cascade ports appears.

Usage Guidelines

The following information appears:

- Port number
- Cascade port
- Ext slot
- Slot identifier for the attached BPE
- Port state
- Link state
- PECSP MAC address
- PE MAC address
- CSPOpen and Loc (local) Rem (remote) status— a value of 1 indicates that a CSP Open request has been sent and acknowledged.
- CSP role flags— when using two CBs with MLAG, the flags indicate which node is performing the role of master (M) and backup (B)

In a redundant CB configuration, the M flag in the output of this command indicates the CB occupying the master role.

Example

The following example shows information for BPEs attached to all cascade ports:

```
# show vpex ports
Port      Cascade Ext  Port  Link  PECSP          PE          CSPOpen
#         Port    Slot State State MAC Address  MAC Address  Loc Rem  Flags
=====
=
1:45      1:45      100   E     A     d8:84:66:88:98:7f d8:84:66:88:98:65 1   1
1:47      1:47      120   E     A     d8:84:66:88:99:14 d8:84:66:88:98:e3 1   1
1:48      1:47      120   E     A     d8:84:66:88:99:16 d8:84:66:88:98:e3 1   1
100:25    100:25     101   E     A     d8:84:66:88:99:1c d8:84:66:88:99:03 1   1
101:26    101:26     122   E     A     d8:84:66:88:99:27 d8:84:66:88:98:f3 1   1
120:50    120:50     121   E     A     d8:84:66:88:98:e8 d8:84:66:88:98:b7 1   1
120:52    120:50     121   E     A     d8:84:66:88:98:e9 d8:84:66:88:98:b7 1   1
122:49    122:50     121   E     A     d8:84:66:88:99:22 d8:84:66:88:98:65 0   0      d
122:50    122:50     121   E     A     d8:84:66:88:99:20 d8:84:66:88:98:e3 0   0      d
=====
=
Port State: D-Disabled, E-Enabled, F-Disabled by link-flap detection,
             L-Disabled due to licensing
Link State: A-Active, R-Ready, NP-Port Not Present, L-Loopback,
Flags: B-CSP Role Backup, C-CSP Role State Complete, d-Dynamic Cascade Port,
       I-CSP Role State Init, M-CSP Role Master, T-CSP Role State
       Transitioning, U-Unknown.
```

History

This command was first available in ExtremeXOS 22.5.

Ring topology support was added in ExtremeXOS 22.7

Platform Availability

This command is available on ExtremeSwitching 5420 and 5520 series switches.

show vpex ports ecp statistics

```
show vpex ports {port_list} ecp statistics
```

Description

Shows bridge port extender (BPE) Edge Control Protocol (ECP) counter information.

Syntax Description

vpex	Specifies Virtual Port Extender (VPEX).
ports	Specifies the cascade port attached to the BPE.
<i>port_list</i>	Specifies the cascade port attached to the BPE.
ecp	Edge Control Protocol.
statistics	Specifies showing ECP counter information.

Default

If you do not specify any ports, information for all cascade ports appears.

Example

The following example shows ECP counter information for BPEs attached to all cascade ports. In this case, there is only one attached BPE on cascade port 1:47:

```
# show vpex ports ecp statistics
Port   RxPkt  RxAck  RxReq  RxSeq  TxPkt  TxAck  TxReq  TxSeq  RxErr  TxErr
=====
1:47   52      0      0     590    52     26     26     26      0      0
=====
Flags   : (*) Active,
          (!) Administratively disabled,
          (A) Auto-calculated timer value
```

History

This command was first available in ExtremeXOS 22.5.

Platform Availability

This command is available on ExtremeSwitching 5420 and 5520 series switches.

show vpex ports statistics

```
show vpex ports {port_list} {statistics} {detail}
```

Description

Shows bridge port extender (BPE) statistics information on an individual per-cascade port basis.

Syntax Description

vpex	Specifies Virtual Port Extender (VPEX).
ports	Specifies the cascade port attached to the BPE.
<i>port_list</i>	Specifies the cascade port attached to the BPE.
statistics	Specifies BPE statistics information on a per-port basis.
detail	Shows detailed BPE statistics information. The detail option shows zero and non-zero counters; when detail is omitted, only non-zero counters appear.

Default

N/A.

Usage Guidelines

This command displays information on an individual port basis (cascade ports can be LAG members). If the **detail** option is not used, only non-zero counts appear.

Example

The following example shows non-detailed BPE statistics for all cascade ports:

```
# show vpex ports statistics

Port: 1:45  BPE: d8:84:66:88:97:fb  Cascade Port: 1:45  Slot: 100  [LAG]
CSP Open: Local: 1 Remote: 1

txRequests      : 9596      rxResponses     : 9596
txResponses     : 0        rxRequests      : 0
rxrqErrors      : 0        rxrspErrors     : 0

Standard Statistics
-----
                        TX Requests                      RX Responses
orgSpecific      : 9596                                9596
                        TX Responses                      RX Requests

Organizationally Specific Statistics
-----
                        TX Requests                      RX Responses
envGet           : 1                                    1
getBulk          : 9594                                9594
portCntrClear    : 1                                    1
                        RX Requests                      TX Responses
No Valid Statistics
=====

Port: 1:46  BPE: d8:84:66:88:97:fc  Cascade Port: 1:45  Slot: 101  [LAG]
CSP Open: Local: 1 Remote: 1

txRequests      : 0        rxResponses     : 0
txResponses     : 0        rxRequests      : 0
rxrqErrors      : 0        rxrspErrors     : 0

Standard Statistics
-----
                        TX Requests                      RX Responses
                        TX Responses                      RX Requests

Organizationally Specific Statistics
-----
                        TX Requests                      RX Responses
No Valid Statistics
                        RX Requests                      TX Responses
No Valid Statistics
```

The following example shows detailed BPE statistics for cascade port 1:45:

```
# show vpex ports 1:45 statistics detail
```

Port: 1:45 BPE: d8:84:66:88:97:fb Cascade Port: 1:45 Slot: 100 [LAG]
 CSP Open: Local: 1 Remote: 1

txRequests	: 9702	rxResponses	: 9702
txResponses	: 0	rxRequests	: 0
rxrqErrors	: 0	rxrspErrors	: 0

Standard Statistics

	TX Requests	RX Responses
cspOpen	: 0	0
extPortCreate	: 0	0
extPortDelete	: 0	0
portParamSet	: 0	0
portParamGet	: 0	0
statusParamSet	: 0	0
echanRegister	: 0	0
echanRegGet	: 0	0
statisticsGet	: 0	0
transitDelaySet	: 0	0
objectGet	: 0	0
objectSet	: 0	0
cnParamGet	: 0	0
cnParamSet	: 0	0
orgSpecific	: 9702	9702
unsupported	: 0	0

	TX Responses	RX Requests
cspOpen	: 0	0
extPortCreate	: 0	0
extPortDelete	: 0	0
portParamSet	: 0	0
portParamGet	: 0	0
statusParamSet	: 0	0
echanRegister	: 0	0
echanRegGet	: 0	0
statisticsGet	: 0	0
transitDelaySet	: 0	0
objectGet	: 0	0
objectSet	: 0	0
cnParamGet	: 0	0
cnParamSet	: 0	0
orgSpecific	: 0	0
unsupported	: 0	0

Organizationally Specific Statistics

	TX Requests	RX Responses
reload	: 0	0
swapImg	: 0	0
imageInfoGet	: 0	0
trstdPrtInfGet	: 0	0
poeParamSet	: 0	0
portCfgSet	: 0	0
portOpStatGet	: 0	0
frmSzeLimSet	: 0	0
frmSzeLimGet	: 0	0
lagBalAlgSet	: 0	0
ecpParamSet	: 0	0
cbMsgRefl	: 0	0
envGet	: 1	1
closeNotif	: 0	0
getBulk	: 9700	9700
localMirrSet	: 0	0

```

debugInfoGet      : 0      0
qCountersGet      : 0      0
qTrafShaperSet    : 0      0
cpuUtilGet        : 0      0
diagSet           : 0      0
addtlCapSet       : 0      0
defPrtQoSParSet   : 0      0
dtaCtlTrfPriSet   : 0      0
ipInterfaceSet    : 0      0
lagSet            : 0      0
portCntrClear     : 1      1
poeFltReasGet     : 0      0
eeeSet            : 0      0
poeGlobParmSet    : 0      0
tftpDownload      : 0      0
sfpNotifSet       : 0      0
diagInfoGet       : 0      0
unsupported        : 0      0

                RX Requests      TX Responses
POE Notif       : 0      0
cbMsgRefl       : 0      0
sfpNotif        : 0      0
unsupported      : 0      0

```

History

This command was first available in ExtremeXOS 22.5.

Platform Availability

This command is available on ExtremeSwitching 5420 and 5520 series switches.

show vpex stacking

```
show vpex stacking
```

Description

In an Extended Edge on stacking environment, this command shows Virtual Port Extender (VPEX) information for all nodes in the topology.

Syntax Description

vpex	Specifies Virtual Port Extender (VPEX).
stacking	Specifies VPEX stacking information.

Default

N/A.

Usage Guidelines

This command is only available in dot1BR controlling bridges and not available on other stackables.

Example

The following example shows VPEX stacking information:

```
Node MAC Address      Slot Flags
-----
*00:11:22:33:44:55   1   eEmc-
00:00:44:33:22:11   2   e--c-
00:99:88:77:66:55   3   e--cA
* - Indicates this node
Flags: (e) VPEX enabled setting is on, (E) VPEX is Active,
       (m) VPEX MLAG mode is configured, (c) VPEX-capable node,
       (A) Active stack node
```

History

This command was first available in ExtremeXOS 31.4.

Platform Availability

This command is available on ExtremeSwitching 5420 and 5520 series switches.

show vpex topology

```
show vpex topology { port port_num } {summary | detail}
```

Description

In an Extended Edge Switching environment, shows the bridge port extender (BPE) topology that is connected to the specified native cascade port.

Syntax Description

vpex	Specifies Virtual Port Extender (VPEX).
port	Specifies topology information for VPEX slots attached to a specified native VPEX cascade port on the controlling bridge (CB).
<i>port_num</i>	Specifies the native VPEX cascade port number.
summary	Includes connected VPEX slot numbers only. Default behavior when port is not specified.
detail	Selects including VPEX slot interconnections. Default behavior when port is specified.

Default

If a port is not specified, then summary (connected VPEX slot numbers only) information appears by default; if a port is specified, then detailed (VPEX slot interconnections) information appears by default.

Usage Guidelines

If a port is not entered or **summary** is selected, all topologies appear (summary mode). Only the native cascade port or MLAG and a list of slots appear in summary mode.

If a port is entered or **detail** is selected, detailed mode information appears. For a ring topology, the command shows both cascades that form a ring, including the dynamically created cascade ports and assigned backup roles. Thus the BPE in each cascade appears twice, one in each cascade.

Severed rings also appear as a pair of cascades with all upstream or cascade ports not being identified as backup. The cascades also indicate that the BPE has been dynamically created (configured) in its position in the cascade. If a native cascade port is not specified, then all topologies appear. Otherwise, only the topology connected to the CB's cascade port appears.

Possible ring states:

- Idle—Ring has been detected, but its configuration has not yet started.
- Init—The new ring is being defined. Dynamic cascade ports required for the redundancy provided by the ring feature are being created.
- Configuring—The ring is being brought up.
- Complete—Ring is fully operational and ready for recovery should a connection be severed.
- Severed—The ring is operating as two cascades.
- Rebalancing—The data plane of a ring is being rebalanced.

For slots configured using virtual MLAG, instead of the native cascade port, the MLAG ID appears.

Example

The following example shows an 8-BPE ring that is formed with two MLAGs. The configured cascades were equal in length. In this case the ring was previously formed, but is now severed. MLAG 10 contains only local native cascade port 1:3. MLAG 11 contains only a native cascade port on the remote MLAG switch. Only slots that exist on the data plane portions of the long cascade remnants after the severance appear:

```
# show vxpe topology port 1:3

Topology: Ring (severed)

Native cascade port 1:3 (MLAG id 10)
  Slot    Upstream  Cascade  Cascade
  ----    -
  108     49         50         109
  109     49         50         110
  ----    -

MLAG peer "vpexpeer" id 11
  Slot    Upstream  Cascade  Cascade
  ----    -
  115     49         50         114
  114     49         50         113
  113     49         50         112
  112     49         50         111
  111     50         49         110
  110     50         40         109
```

```

-----
* - Port is a member of a load sharing group
Flags: (b) Backup Upstream Port, (B) Backup Cascade Port,
       (c) Common ring link on this BPE
       (d) Dynamically Created Cascade Port,
       (E) Empty slot, (S) Severed ring link on this BPE.

```

The following example is a summary display of all topologies:

```

# show vpex topology

Topology: Ring (Complete)
  Native cascade port 1:1
    Slots: 100, 101, 102, 103, 104, 105, 106, 107
  Native cascade port 1:2
    Slots: 107, 106, 105, 104, 103, 102, 101, 100

Topology: Ring (Severed)
  Native cascade port 1:3 (MLAG id 10)
    Slots: 108, 109
  MLAG ID 11
    Slots: 115, 114, 113, 112, 111, 110

Topology: Cascade
  Native cascade port 1:4
    Slots: 116, 117, 118, 119, 120

```

History

This command was first available in ExtremeXOS 22.7.

Platform Availability

This command is available on ExtremeSwitching 5420 and 5520 series switches.

show vpls

```
show vpls {{vpls_name}} {peer ipaddress} {detail} | summary
```



Note

This command has been replaced with the following command: `show l2vpn {vpls {{vpls_name}} | vpws {{vpws_name}} {peeripaddress} {detail} | summary}`. This command is still supported for backward compatibility, but it will be removed from a future release, so we recommend that you start using the new command.

Description

Displays VPLS and H-VPLS configuration and status information.

Syntax Description

<i>vpls_name</i>	Displays information for the specified vpls.
<i>ipaddress</i>	Specifies a VPLS peer for which to display information.
detail	Displays additional information in comprehensive detail format.

Default

N/A.

Usage Guidelines

The `show vpls` command (without any optional parameters) displays all currently configured VPLS instances for the switch. The summarized list of VPLS instances is displayed in alphabetical order based on the *vpls_name*. Peers are displayed in the reverse of the order they were added.

When you specify a VPLS peer, the display includes a list of all PWs established to the peer, the PW status and PW ID, and information about each VPLS to which this peer belongs.

The `show 12vpn` command describes the display fields that appear when this command is entered with the **detail** option.

Table 45: Selected show vpls Field Definitions

Field	Definition
VPLS Name	VPLS instance or domain name.
VPN ID	Virtual Private Network identifier.
Source Address	Source IP address.
VCCV Status	Virtual Circuit Connectivity Verification (VCCV) feature status, which is either Enabled or Disabled .
VCCV Interval Time	Displays the configured VCCV interval time.
VCCV Fault Multiplier	Displays the configured VCCV fault multiplier.
Redundancy Type	Displays the configured VPLS redundancy type, which is EAPS , ESRP , or None .
Service Interface	Displays a <i>VLAN</i> or <i>VMAN</i> interface name.
Admin State	Displays the administrative state of the VPLS, which is either Enabled or Disabled .
Oper State	Displays the operational state of the VPLS, which is either Enabled or Disabled .
MTU	Displays the maximum transmission unit (MTU) size for the VPLS.
Ethertype	Displays the ethertype for the service interface.
.1q tag	Displays the 802.1q priority tag for the VPLS.
Peer IP	Displays the IP address for the VPLS peer.

Table 45: Selected show vpls Field Definitions (continued)

Field	Definition
PW State	<p>PW State represents the state, or status, of a PW. The possible PW state values are:</p> <ul style="list-style-type: none"> • UP—The PW is fully operational and installed in hardware. Traffic is forwarded over PW and VPLS service VLAN/VMAN. • Down—The PW is not operational and is not installed in hardware. This only happens when the VPLS instance is disabled, VPLS service is disabled, or there is no service VLAN assigned to the VPLS. No traffic is forwarded. • Sgnl—The PW is in a signalling state. The PW is not operational, and no traffic is forwarded. This can occur for a number of reasons, including: No LDP adjacency to peer, No transport LSP to peer, No VC LSP to peer. • Remote peer not configuredReady—The PW has been signalled, but it has not been installed in hardware. Traffic is not forwarded. The PW can be in a Ready state for a number of reasons, including: <ul style="list-style-type: none"> ◦ The VPLS instance is configured for <i>EAPS</i> redundancy, and the EAPS shared port associated with this VPLS instance is Connected. ◦ The VPLS instance is configured for <i>ESRP</i> redundancy, and the ESRP domain associated with this VPLS instance is Slave. ◦ The service VLAN associated with this VPLS instance is down. ◦ The remote peer has signalled that it has a fault (remote PW status). The remote peer may have a fault due to it's service VLAN being down.
PW Uptime	PW Uptime is the elapsed time that the PW has been in the UP state.
PW Installed	PW Installed is a flag to indicate whether the PW is installed in hardware or not. If the PW is in the UP state, this field is True, otherwise, this field is False.
Local PW Status	<p>Local PW Status displays the VC status of the local PW. The values are:</p> <ul style="list-style-type: none"> • No Faults—No faults detected. • PW-Tx—Local PSN-facing PW transmit fault. This is set if there is a problem with the VPLS transport LSP. • PW-Rx—Local PSN-facing PW receive fault. This is set if there is a problem with the VPLS transport LSP. • Att-Tx—Local attachment circuit transmit fault. This is set if there is a problem with the VPLS service VLAN. • Att-Rx—Local attachment circuit receive fault. This is set if there is a problem with the VPLS service VLAN. • Not Forwarding—The local PW is not forwarding. Look for more information in the PW State field. For example, if VPLS is configured for EAPS redundancy, the Local PW Status is Not Forwarding and the PW State is Ready whenever the EAPS Shared Port state is Connected.
Remote PW Status	Remote PW Status is the VC status of the remote PW. The values for this field are the same values as for Local PW Status.

Table 45: Selected show vpls Field Definitions (continued)

Field	Definition
PW Mode	PW Mode describes how the PW was configured. The values are: <ul style="list-style-type: none"> Core-to-Core—This VPLS instance is a core node, and the other end of the PW connects to a core node. Core-to-Spoke—This VPLS instance is a core node, and the other end of the PW connects to a spoke node. This is for HVPLS. Spoke-to-Core—This VPLS instance is a spoke node, and the other end of the PW connects to a core node. This is for HVPLS.
Transport LSP	Transport LSP is the LSP that is used to forward frames over the PW. When an LDP LSP is used as a transport, the display shows LDP LSP (Not configured). If an RSVP LSP is used, the name of the RSVP LSP being used as a transport LSP is displayed. An RSVP LSP can be specified as the LSP to use during VPLS configuration.
Next Hop I/F	Displays the interface name for the next hop router.
Next Hop Addr	Displays the interface IP address for the next hop router.
PW Rx Label	Receive label for the VPLS PW.
PW Rx Pkts	Total packets received on the VPLS PW.
PW Rx Bytes	Total bytes received on the VPLS PW.
Tx Label	Transmit label for the LSP.
PW Tx Label	Transmit label for the VPLS PW.
PW Tx Pkts	Total packets transmitted on the VPLS PW.
PW Tx Bytes	Total bytes transmitted on the VPLS PW.

Example

The following shows the display that appears when you enter the command without the **detail** option:

```
Switch.38 # show vpls
L2VPN Name      VPN ID  Flags  Services Name  Peer IP      State  Flags
-----
Pws-3344        20      EAX--W  NONE           11.100.100.219 Up      C--NV-
jwcvpls         99      EAX--L  torix          11.100.100.219 Up      C--NV-
keeper          90      EAX--L  NONE           11.100.100.219 Up      ----V-
pws-1           2009    EAX--W  pwserve        11.100.100.219 Up      ----V-
pws-10          70      EAX--W  NONE           11.100.100.219 Up      ---NV-
pws-2           2008    EAX--W  pw2serve       11.100.100.219 Up      ---NV-
pws-3           2007    EAX--W  NONE           11.100.100.219 Up      ---NV-
sarsparilla     80      EAX--W  NONE           11.100.100.219 Down    C--N--
whoopwoo        100     EAX--L  NONE           11.100.100.219 Down    C--N--
VPN Flags: (E) Admin Enabled, (A) Oper Active, (I) Include Tag,
(X) Exclude Tag, (T) Ethertype Configured,
(V) VCCV HC Enabled, (W) VPN Type VPWS, (L) VPN Type VPLS
Peer Flags: (C) Core Peer, (S) Spoke Peer, (A) Active Core,
(p) Configured Primary Core, (s) Configured Secondary Core,
(N) Named LSP Configured, (V) VCCV HC Capabilities Negotiated,
(F) VCCV HC Failed
-----
Total number of configured L2VPNs: 9
```

```

Total number of active L2VPNs:      3
Total number of configured PWs:     4
Total number of active PWs:        3
PWs auto-selecting transport LSP:   1
PWs configured with a transport LSP: 3
PWs using LDP for transport:        0
PWs using RSVP for transport:       4
PWs using static for transport:     0

```

The following shows summary L2 VPN information for the specified VPLS peer:

```

Switch.451 # sh vpls peer 2.2.2.2
L2VPN Name      VPN ID Flags  Services Name  Peer IP      State Flags
-----
vsl             105     EAX--L cust1   2.2.2.2      UP          CAp-V-
VPN Flags: (E) Admin Enabled, (A) Oper Active, (I) Include Tag,
(X) Exclude Tag, (T) Ethertype Configured,
(V) VCCV HC Enabled, (W) VPN Type VPWS, (L) VPN Type VPLS
Peer Flags: (C) Core Peer, (S) Spoke Peer, (A) Active Core,
(p) Configured Primary Core, (s) Configured Secondary Core,
(N) Named LSP Configured, (V) VCCV HC Capabilities Negotiated,
(F) VCCV HC Failed

```

The following shows detailed L2 VPN information for the specified VPLS peer:

```

Switch.452 # sh vpls peer 11.100.100.210 detail
VPLS Name : vpls10
VPN ID      : 10
Source Address : 11.100.100.212
VCCV Status  : Disabled
VCCV Interval Time : 5 sec.
VCCV Fault Multiplier : 4
L2VPN Type   : VPLS
Service Interface : vlan10
Peer IP : 11.100.100.210
PW State      : Up
PW Uptime     : 18d:0h:28m:26s
PW Installed  : True
Local PW Status : No Faults
Remote PW Status : No Faults
PW Mode       : Core-to-Core
Transport LSP  : LDP LSP (Not Configured)
Next Hop I/F   : o6vlan1
Next Hop Addr  : 12.182.0.216
Tx Label      : 0x00010
PW Rx Label    : 0x80405
PW Rx Pkts     : 3806161633
PW Rx Bytes    : 912385942
MAC Limit     : No Limit
VCCV HC Status : Not Sending (VCCV Not Enabled For This VPLS)
CC Type       : Rtr Alert
CV Type       : LSP Ping
Send Next Pkt : --
Total Failures : 0
Last Failure Tm : --

```

The following command shows the status of L2VPN Sharing configuration. If L2VPN Sharing is enabled, and more than 1 Transport LSP is configured, the output will display the status of each Transport LSP. The following example shows VPLS instance “vpls1” with LSP sharing enabled, with 2 Transport LSPs configured and programmed into HW.

```

* dut1.11 # show vpls detail
L2VPN Name: vpls1
VPN ID      : 91
Admin State : Enabled

```

```

Source Address      : 20.20.20.91      Oper State   : Enabled
VCCV Status        : Disabled          MTU          : 1500
VCCV Interval Time : 5 sec.            Ethertype    : 0x8100
VCCV Fault Multiplier : 4              .1q tag      : exclude
L2VPN Type         : VPLS              Redundancy   : None
Service Interface  : vlan91

```

```

Peer IP: 20.20.20.83
PW State           : Up
PW Uptime          : 2d:22h:28m:44s
PW Installed       : True
Local PW Status    : No Faults
Remote PW Status   : No Faults
Remote I/F MTU     : 1500
PW Mode            : Core-to-Core
Transport LSP      : lsp1 (Configured)
  Next Hop I/F     : vlan21
  Next Hop Addr    : 10.0.21.92        Tx Label   : 0x00074
Transport LSP      : lsp2 (Configured)
  Next Hop I/F     : vlan21
  Next Hop Addr    : 10.0.21.92        Tx Label   : 0x00077
PW Rx Label        : 0x00075          PW Tx Label  : 0x00085
PW Rx Pkts         : 0                PW Tx Pkts   : 0
PW Rx Bytes        : 0                PW Tx Bytes  : 0
MAC Limit          : No Limit
VCCV HC Status     : Not Sending (VCCV Not Enabled For This L2VPN)
  CC Type          : Rtr Alert          Total Pkts Sent : 0
  CV Type          : LSP Ping           Total Pkts Rcvd : 0
  Send Next Pkt    : --
  Total Failures   : 0                  Pkts During Last Failure : 0
  Last Failure Tm  : --

```

```

-----
Total number of configured L2VPNs: 1
Total number of active L2VPNs: 1
Total number of configured PWs: 1
Total number of active PWs: 1
Total number of ready PWs: 0
PWs auto-selecting transport LSP: 0
PWs configured with a transport LSP: 1
PWs using LDP for transport: 0
PWs using RSVP for transport: 1
PWs using static for transport: 0

```

* - Indicates LSP is configured, but not installed in HW because L2VPN Sharing is disabled.

The following example shows VPLS instance “vp1s1” with LSP sharing disabled, with 2 Transport LSPs configured, but only 1 LSP programmed into HW (because sharing is disabled).

```

* dut1.11 # show vp1s detail
L2VPN Name: vp1s1
  VPN ID      : 91      Admin State   : Enabled
Source Address : 20.20.20.91 Oper State   : Enabled
VCCV Status    : Disabled MTU          : 1500
VCCV Interval Time : 5 sec. Ethertype    : 0x8100
VCCV Fault Multiplier : 4    .1q tag      : exclude
L2VPN Type     : VPLS      Redundancy   : None
Service Interface : vlan91

Peer IP: 20.20.20.83
  PW State      : Up
  PW Uptime     : 2d:22h:28m:44s
  PW Installed   : True

```



```

Local PW Status      : No Faults
Remote PW Status     : No Faults
Remote I/F MTU       : 1500
PW Mode              : Core-to-Core
Transport LSP        : lsp1 (Configured)
  Next Hop I/F       : vlan21
  Next Hop Addr      : 10.0.21.92          Tx Label   : 0x00074
Transport LSP      : lsp2 (Configured*)
  Next Hop I/F     : vlan21
  Next Hop Addr    : 10.0.21.92          Tx Label    : 0x00077
PW Rx Label          : 0x00075           PW Tx Label   : 0x00085
PW Rx Pkts           : 0                 PW Tx Pkts    : 0
PW Rx Bytes          : 0                 PW Tx Bytes    : 0
MAC Limit            : No Limit
VCCV HC Status       : Not Sending (VCCV Not Enabled For This L2VPN)
  CC Type             : Rtr Alert         Total Pkts Sent : 0
  CV Type             : LSP Ping          Total Pkts Rcvd : 0
  Send Next Pkt       : --
  Total Failures      : 0                 Pkts During Last Failure : 0
  Last Failure Tm     : --

```

```

-----
Total number of configured L2VPNs:    1
Total number of active L2VPNs:        1
Total number of configured PWs:        1
Total number of active PWs:            1
Total number of ready PWs:             0
PWs auto-selecting transport LSP:      0
PWs configured with a transport LSP:   1
PWs using LDP for transport:            0
PWs using RSVP for transport:          1
PWs using static for transport:         0

```

* - Indicates LSP is configured, but not installed in HW because L2VPN Sharing is disabled.

This following shows VPLS instance “vpls1” with LSP sharing enabled, 2 LSPs configured, but only 1 LSP programmed into HW. (LSP2 is disabled).

```

* dut1.13 # show vpls detail      L2VPN
Name: vpls1
  VPN ID                : 91                Admin State   : Enabled
  Source Address        : 20.20.20.91       Oper State    : Enabled
  VCCV Status           : Disabled           MTU          : 1500
  VCCV Interval Time    : 5 sec.            Ethertype     : 0x8100
  VCCV Fault Multiplier : 4                 .1q tag       : exclude
  L2VPN Type            : VPLS               Redundancy    : None
  Service Interface     : vlan91

Peer IP: 20.20.20.83      PW
  State                 : Up
  PW Uptime              : 2d:22h:30m:13s
  PW Installed           : True
  Local PW Status        : No Faults
  Remote PW Status       : No Faults
  Remote I/F MTU         : 1500
  PW Mode                : Core-to-Core
  Transport LSP          : lsp1 (Configured)
    Next Hop I/F         : vlan21
    Next Hop Addr        : 10.0.21.92          Tx Label     : 0x00074
  Transport LSP          : lsp2 (Configured)
    Next Hop I/F         : --
    Next Hop Addr        : --                  Tx Label     : --
  PW Rx Label            : 0x00075           PW Tx Label   : 0x00085

```

```

PW Rx Pkts      : 0                      PW Tx Pkts    : 0
PW Rx Bytes     : 0                      PW Tx Bytes   : 0
MAC Limit       : No Limit
VCCV HC Status  : Not Sending (VCCV Not Enabled For This L2VPN)
  CC Type       : Rtr Alert                Total Pkts Sent : 0
  CV Type       : LSP Ping                 Total Pkts Rcvd : 0
  Send Next Pkt : --
  Total Failures : 0                      Pkts During Last Failure : 0
  Last Failure Tm: --

-----

Total number of configured L2VPNs: 1
Total number of active L2VPNs: 1
Total number of configured PWs: 1
Total number of active PWs: 1
Total number of ready PWs: 0
PWs auto-selecting transport LSP: 0
PWs configured with a transport LSP: 1
PWs using LDP for transport: 0
PWs using RSVP for transport: 1
PWs using static for transport: 0

* - Indicates LSP is configured, but not installed in HW because L2VPN Sharing is
disabled.

```

The following shows VPLS instance “vpls1” with LSP Sharing enabled, with 5 Transport LSPs configured and programmed into HW.

```

* dut1.19 # show vpls detail
L2VPN Name: vpls1
VPN ID      : 91                      Admin State : Enabled
Source Address : 20.20.20.91          Oper State  : Enabled
VCCV Status   : Disabled              MTU         : 1500
VCCV Interval Time : 5 sec.           Ethertype   : 0x8100
VCCV Fault Multiplier : 4              .1q tag     : exclude
L2VPN Type    : VPLS                  Redundancy  : None
Service Interface : vlan91

Peer IP: 20.20.20.83
  PW State      : Up
  PW Uptime     : 2d:22h:31m:54s
  PW Installed  : True
  Local PW Status : No Faults
  Remote PW Status : No Faults
  Remote I/F MTU : 1500
  PW Mode       : Core-to-Core
  Transport LSP : lsp1 (Configured)
    Next Hop I/F : vlan21
    Next Hop Addr : 10.0.21.92          Tx Label    : 0x00074
  Transport LSP : lsp2 (Configured)
    Next Hop I/F : vlan21
    Next Hop Addr : 10.0.21.92          Tx Label    : 0x00077
  Transport LSP : lsp3 (Configured)
    Next Hop I/F : vlan21
    Next Hop Addr : 10.0.21.92          Tx Label    : 0x00075
  Transport LSP : lsp4 (Configured)
    Next Hop I/F : vlan21
    Next Hop Addr : 10.0.21.92          Tx Label    : 0x00076
  Transport LSP : lsp5 (Configured)
    Next Hop I/F : vlan21
    Next Hop Addr : 10.0.21.92          Tx Label    : 0x00078

```

```

PW Rx Label      : 0x00075          PW Tx Label   : 0x00085
PW Rx Pkts       : 0                PW Tx Pkts    : 0
PW Rx Bytes      : 0                PW Tx Bytes   : 0
MAC Limit        : No Limit
VCCV HC Status   : Not Sending (VCCV Not Enabled For This L2VPN)
  CC Type        : Rtr Alert          Total Pkts Sent : 0
  CV Type        : LSP Ping           Total Pkts Rcvd : 0
  Send Next Pkt  : --
  Total Failures : 0                  Pkts During Last Failure : 0
  Last Failure Tm: --

```

```

-----
Total number of configured L2VPNs: 1
Total number of active L2VPNs: 1
Total number of configured PWs: 1
Total number of active PWs: 1
Total number of ready PWs: 0
PWs auto-selecting transport LSP: 0
PWs configured with a transport LSP: 1
PWs using LDP for transport: 0
PWs using RSVP for transport: 1
PWs using static for transport: 0

```

* - Indicates LSP is configured, but not installed in HW because L2VPN Sharing is disabled.

History

This command was first available in ExtremeXOS 11.6.

This command was updated to display flags for H-VPLS spoke nodes and protected VPLS and H-VPLS in ExtremeXOS 12.1.

The output for this command was modified in ExtremeXOS 12.2.2.

The output for this command was modified in ExtremeXOS 15.4 to display LSP sharing information.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

show vpls peer l2pt

```

show {l2vpn} [vpls vpls_name | vpws] vpws_name] {peer ipaddress} l2pt
{detail}

```

Description

Displays the VPLS peer L2PT configuration and status of a service.

Syntax Description

l2vpn	Specifies the Layer 2 Virtual Private Network.
vp1s	Specifies the Virtual Private LAN Service over MPLS.
<i>vp1s_name</i>	Specifies the VPLS VPN.
vpws	Specifies the Virtual Private Wire Service over MPLS.
<i>vpws_name</i>	Specifies the VPWS VPN.
peer	Specifies the peer.
<i>ipaddress</i>	Specifies the IPv4 address.
l2pt	Specifies Layer 2 protocol configuration.
detail	Specifies to show L2PT configuration and status in detail.

Default

Disabled.

Usage Guidelines

Use this command to display the L2PT configuration and status of a service.

Example

The following is an example of the `show vp1s peer l2pt` command:

```
# show l2vpn vp1s cust1 l2pt

Interface      L2PT Profile Name
-----
123.112.111.111 my_other_l2pt_prof
```

History

This command was first available in ExtremeXOS 15.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show vrrp

```
show vrrp {virtual-router {vr_name}} {detail}
```

Description

Displays *VRRP* configuration information for all VRRP VLANs.

Syntax Description

<i>vr_name</i>	Specifies a virtual router (VR) for which to display VRRP information.
detail	Specifies more detailed VRRP information.

Default

N/A.

Usage Guidelines

The following table describes the significant fields and values that can appear when you enter the different forms of this command:

Field	Description
Advertisement Interval	Indicates the configured advertisement interval.
Authentication	If configured, identifies the VRRP simple password.
FR	Shows if Fabric Routing is enabled.
HM	Shows if Host Mobility is enabled.
IPv6 Router Advertisement Mask	IPv6 router advertisement mask.
Master Mac Address	The MAC address of the master VRRP router.
P	Indicates that the VRRP <i>VLAN</i> is configured for preempt mode, which controls whether a higher priority backup router preempts a lower priority master.
Preempt	Indicates that the VRRP <i>VLAN</i> is configured for preempt mode, which controls whether a higher priority backup router preempts a lower priority master.
Pri Priority	The priority value of the VRRP <i>VLAN</i> .
State	The current state of the VRRP router. The state includes the following: Init —The VRRP router is in the initial state. Backup—The VRRP router is a backup router. Master—The VRRP router is the master router.
T	Indicates the configured advertisement interval.
TP	Indicates the number of tracked pings.
TR	Indicates the number of tracked routes.
Tracked IP Routes	If configured, displays the IP address and subnet mask of the tracked route(s).
Tracked Pings	If configured, displays the: Target IP address you are pinging. Number of seconds between pings to the target IP address. Number of misses allowed before this entry is considered to be failing.
Tracked VLANs	If configured, displays the name of the tracked <i>VLAN</i> (s).
Tracking Mode	Indicates the VRRP tracking mode, which is either ALL or ANY.
TV	Indicates the number of tracked <i>VLAN</i> s.

Field	Description
Virtual IP Addr Virtual IP Addresses	If configured, the virtual IPv4 or IPv6 address associated with the VRRP VLAN.
Virtual Link-Local Address	Virtual IPv6 link local address configured on the interface.
VLAN	The name of the VRRP VLAN.
VLAN Name	The name of the VRRP VLAN and whether VRRP is enabled or disabled on the VLAN. The enable/disable state appears as follows: En—VRRP is enabled on this VLAN.Ds—VRRP is disabled on this VLAN.
VRID	The VRRP Router Identification number for the VRRP instance.
VRRP	The enabled/disabled state of VRRP on the VLAN.

Example

The following shows VRRP information for the current VR context:

```
# show vrrp

          Virtual          Master
VLAN Name VRID Pri IP Address      State MAC Address      TP/TR/TV/P/T    /FR/G/HM
tenant(En) 0001 255 10.1.1.1      INIT 00:00:5e:00:01:01  0 0 0 Y 1      N N N
tenant(Ds) 0002 100 NONE          INIT NONE          0 0 0 Y 1      N N Y

En-Enabled, Ds-Disabled, Pri-Priority, T-Advert Timer, P-Preempt
TP-Tracked Pings, TR-Tracked Routes, TV-Tracked VLANs, FR-Fabric Routing,
G-Group, HM-Host Mobility
```

The following example displays detailed VRRP status information for the current VR context:

```
# show vrrp detail
VLAN: vlan23   VRID: 1       VRRP: Disabled State: INIT
Virtual Router: VR-Default
Priority: 100(backup) Advertisement Interval: 1 sec
Version: v3-v2 Preempt: Yes  Preempt Delay: 0 sec
Virtual IP Addresses:
Accept mode: Off          Group : ExtremeNet(Enabled)
Host-Mobility: On
Host-Mobility Excluded-Ports:
Checksum: Include pseudo-header
Tracking mode: ALL
Tracked Pings: -
Tracked IP Routes: -
Tracked VLANs: -
Fabric Routing: Off
```

History

This command was first available in ExtremeXOS 10.1.

Support for virtual routers was added in ExtremeXOS 12.0.

Support for IPv6 was added in ExtremeXOS 12.7.

Support for VRRP groups added in ExtremeXOS 22.2.

Host mobility information added in ExtremeXOS 22.4.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the VRRP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

show vrrp group

```
show vrrp group group_name
```

Description

This command shows VRRP group status.

Syntax Description

group	Form a group of VRRP VRs to operate in high-scale mode.
<i>group_name</i>	Specifies the VRRP group name.

Default

None.

Example

The following example shows VRRP group information for group "ExtremeNet":

```
show vrrp group ExtremeNet
Group Name: ExtremeNet      Group Status : Enabled

      Virtual      Master
VLAN Name VRID Pri IP Address      State MAC Address      TP/TR/TV/P/T      /FR
Primary-VR:
  v1 (En) 0001 225 100.0.0.100      MSTR 00:00:5e:00:01:01 0 0 0 Y 1      Y
Secondary-VR(s):
  v10 (En) 0001 225 100.10.0.100      MSTR 00:00:5e:00:01:01 0 0 0 N 40      Y
  v11 (En) 0001 200 100.11.0.100      MSTR 00:00:5e:00:01:01 0 0 0 N 40      Y

En-Enabled, Ds-Disabled, Pri-Priority, T-Advert Timer, P-Preempt,
TP-Tracked Pings, TR-Tracked Routes, TV-Tracked VLANs, FR- Fabric Routing

Total number of member VRs of group : 3
```

History

This command was first available in ExtremeXOS 22.2.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the VRRP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

show vrrp vlan

```
show vrrp [ {vlan} vlan_name | vlan vlan_list] {stats}
```

Description

Displays VRRP information for a particular VLAN.

Syntax Description

<i>vlan_name</i>	Specifies the VRRP VLAN.
<i>vlan_list</i>	Specifies the VRRP VLAN list of IDs.
stats	Specifies statistics for a particular VLAN.

Default

N/A.

Usage Guidelines

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

If you specify the command without the stats keyword, the following VRRP information appears:

- VLAN—The name of the VRRP VLAN.
- VRID—The VRRP Router Identification number for the VRRP instance.
- VRRP—The enabled/disabled state of VRRP on the VLAN.
- State—The current state of the VRRP router. The state includes the following:
 - Init—The VRRP router is in the initial state.
 - Backup—The VRRP router is a backup router.
 - Master—The VRRP router is the master router.
- Priority—The priority value of the VRRP VLAN.
- Advertisement Interval—Indicates the configured advertisement interval.
- Preempt—Indicates that the VRRP VLAN is configured for preempt mode, which controls whether a higher priority backup router preempts a lower priority master.
- Fabric-routing—If configured IP forwarding capability for VRRP Backup routers is enabled
- Host-mobility—Exportable Host Route learning via ARP/ND on the specified VLAN and VRID
- Host-Mobility Exclude-Ports—If configured these ports will not generate host-route
- Checksum—Include or exclude pseudo header for IPv4 address family

- Authentication—If configured, identifies the VRRP simple password.
- Virtual IP Addresses—If configured, the virtual IP address associated with the VRRP VLAN.
- Tracked Pings—If configured, displays the:
 - Target IP address you are pinging.
 - Number of seconds between pings to the target IP address.
 - Number of misses allowed before this entry is considered to be failing.
- Tracked IP Routes—If configured, displays the IP address and subnet mask of the tracked route(s).
- Tracked VLANs—If configured, displays the name of the tracked VLAN(s).

If you specify the stats keyword, you see counter and statistics information for the specified VRRP VLAN.

Example

The following example displays configuration information for the specified VRRP VLAN:

```
# show vrrp vlan v1
VLAN: v1          VRID: 1          VRRP: Disabled State: INIT
Virtual Router: VR-Default
Priority: 100(backup) Advertisement Interval: 1 sec
Version: v2       Preempt: Yes     Preempt Delay: 0 sec
Authentication: simple-password key: none
Virtual IP Addresses:
Accept mode: Off
Host-Mobility: Off
Host-Mobility Exclude-Ports:
Tracking mode: ALL
Tracked Pings: -
Tracked IP Routes: -
Tracked VLANs: -
Fabric Routing: Off

* indicates a tracking condition has failed
```

The following example displays statistics for VLAN vrrp-1:

```
#show vrrp vlan v1 stats

VLAN v1, VR ID 1

VRID Errors           : 0   Authentication Type Mismatch : 0
Version Errors        : 0   Invalid Authentication Type   : 0
Checksum Errors       : 0   Authentication Failures      : 0
TTL Errors            : 0   Advertisement Interval Errors : 0
Packet Length Errors  : 0   Address List Errors           : 0
Advertisements GotV3 DropV2 : 0
Advertisements Accepted : 0   Priority Zero Packets Received : 0
Become Master         : 0   Priority Zero Packets Sent      : 0
```

History

This command was first available in ExtremeXOS 10.1.

The `vlan_list` variable was added in ExtremeXOS 16.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the VRRP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

show wredprofile

```
show wredprofile {ports [port_list | all]}
```

Description

Displays WRED configuration data for the specified ports or all ports.

Syntax Description

<i>port_list</i>	Specifies a list of slots and ports to display. Specify ports in the following formats: 3-5, 2:5, 2:6-2:8.
all	Specifies that this command applies to all ports on the device.

Default

N/A.

Usage Guidelines

If no port or port list is specified, this command displays the default WRED configuration values.

Example

The following example displays the WRED settings for port 2:

```
# show wredprofile ports 2

Port: 2
WRED configuration parameters
=====
QoS      Packet      Min      Max      Max      Avg.
Profile  Type        Color    Thresh   Thresh   Drop-Rate  Weight  ECN
=====
  QP1    TCP         Green    100%    100%    100%      4%     Off
  QP1    TCP         Red      100%    100%    100%      4%     Off
  QP1    non-TCP    Any      100%    100%    100%      4%     Off
  QP1    non-TCP    Red      100%    100%    100%      4%     Off

  QP5    TCP         Green    100%    100%    100%      5%     On
  QP5    TCP         Red      100%    100%    100%      5%     On
  QP5    non-TCP    Any      100%    100%    100%      5%     On
  QP5    non-TCP    Red      100%    100%    100%      5%     On

  QP8    TCP         Green    100%    100%    100%      4%     Off
```

History

This command was first available in ExtremeXOS 12.7.

ECN column added in ExtremeXOS 22.3.

Platform Availability

This command is available on the ExtremeSwitching 5320 and 5520 series switches.

show xml-notification configuration

```
show xml-notification configuration {target}
```

Description

Displays the configuration of the Web server target.

Syntax Description

<i>target</i>	Specifies an alpha numeric string that identifies the configured target.
---------------	--

Default

N/A.

Usage Guidelines

Use this command to display information about the configuration of the Web server target. If a target is not specified, all configured targets are displayed.

Example

The following command displays the configuration of the configured targets:

```
show xml-notification configuration
```

The following is sample output from this command:

```
Target Name           : sqa
**Server URL          : http://10.255.129.22:8080/xos/webservice (VR-Mgmt)
Server User Name      : admin
Enabled               : yes
Queue Size            : 100
Connection Status     : connected
Configured Modules    : ems,idmgr
Target Name           : epi
**Server URL          : http://10.255.59.6:8080/xos/webservice (VR-Finance)
Server User Name      : admin
Enabled               : yes
Queue Size            : 100
```

```

Connection Status      : connected
Configured Modules     : ems
Target Name            : test3
**Server URL           : https://10.120.91.64:8443/xos/webservice (VR-Mgmt)
Server User Name       : admin
Enabled                : yes
Queue Size             : 100
Connection Status      : not connected
Configured Modules     : ems
Target Name            : testingcorrect
**Server URL           : http://10.66.254.211:8080/xos/webservice (VR-Mgmt)
Server User Name       : admin
Enabled                : no
Queue Size             : 100
Connection Status      : not connected
Configured Modules     : idMgr,ems

```



Note

When a particular VR has been specified in the configuration process, that VR is displayed next to the URL. When no VR is specified since the parameter is optional, the default VR supplied by the XML client is VR-Mgmt. When you are using a version released before the virtual router option was added, VR-Mgmt is displayed.

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

show xml-notification statistics

```
show xml-notification statistics {target}
```

Description

Displays statistics for of the Web server target.

Syntax Description

<i>target</i>	Specifies an alpha numeric string that identifies the configured target.
---------------	--

Default

N/A.

Usage Guidelines

Use this command to display the connection status, enable status and event statistics of the Web server target. If a target is not specified, all configured targets are displayed.

Example

The following command displays statistics for all of the configured targets:

```
show xml-notification statistics
```

The following is sample output from this command:

```
Target Name           : epi
Server URL            : http://10.255.129.22:8080/xos/webservice
Server Queue Size     : 100
Enabled               : yes
Connection Status     : connected
Events Received       : 450
Connection Failures   : 0
Events Sent Success   : 450
Events Sent Fail      : 0
Events Dropped        : 0
Target Name           : epi
Server URL            : http://10.255.59.6:8080/xos/webservice
Server Queue Size     : 100
Enabled               : yes
Connection Status     : fail
Events Received       : 31
Connection Failures   : 3
Events Sent Success   : 2
Events Sent Fail      : 29
Events Dropped        : 0
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

ssh2

```
ssh2 {cipher cipher} {mac mac} {port portnum} {compression [on | off]}
    {user username} {username} [host | ipaddress] {remote command} {vr
vr_name}
```

Description

Initiates an SSH2 client session to a remote SSH2 server.

Syntax Description

<i>cipher</i>	Specifies the name of the cipher. Possible values are: <ul style="list-style-type: none"> • aes128-cbc • aes128-ctr • aes192-cbc • aes192-ctr • aes256-cbc • aes256-ctr • chacha20-poly1305@openssh.com • rijndael-cbc@lysator.liu.se
<i>mac</i>	Specifies the name of the Message Authentication Code. Possible values are: <ul style="list-style-type: none"> • hmac-md5 • hmac-md5-96 • hmac-md5-96-etm@openssh.com • hmac-md5-etm@openssh.com • hmac-sha1 • hmac-sha1-96 • hmac-sha1-96-etm@openssh.com • hmac-sha1-etm@openssh.com • hmac-sha2-256 • hmac-sha2-256-etm@openssh.com • hmac-sha2-512 • hmac-sha2-512-etm@openssh.com
<i>portnum</i>	Specifies the TCP port number to be used for communicating with the SSH2 client. The default is port 22.
on	Specifies that the data is to be compressed.
off	Specifies that compression is not to be used. This is the default.
<i>username</i>	Specifies a login name for the remote host, as an alternate to the username@host parameter. Can be omitted if it is the same as the username on the switch.
<i>host</i>	Specifies the name of the remote host.
<i>ipaddress</i>	Specifies the IP address of the remote host.
<i>remote command</i>	Specifies a command to be passed to the remote system for execution. The switch does not support remote commands. The option is only valid if the remote system is a system, such as a UNIX workstation, that accepts remote commands.
<i>vr_name</i>	Specifies the virtual router. The default virtual router is <i>VR-Mgmt</i> . Note: User-created VRs are supported only on the platforms listed for this feature in the Switch Engine 32.3 Feature License Requirements document.

Default

The default settings for SSH2 parameters are as follows:

- cipher—the full cipher list
- mac—the full Message Authentication Code list
- port—22
- compression—off
- vr_name—VR-Mgmt

Usage Guidelines

SSH2 does not need to be enabled on the switch in order to use this command.

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

Host Name, User Name, and Remote IP Address Character Restrictions

This section provides information about the characters supported by the switch for host names and remote IP addresses.

When specifying a host name, user name, or remote IP address, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-) Permitted for host and user names
- Underscore (_) Permitted for host and user names
- Colon (:) Permitted for host names and remote IP addresses
- At symbol (@) Permitted only for user names

When naming the host, creating a user name, or configuring the IP address, remember the requirements listed above.

Remote Filename Character Restrictions

This section provides information about the characters supported by the switch for remote filenames.

When specifying a remote filename, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-)

- Underscore (_)
- Slash (/)

When naming a remote file, remember the requirements listed above.

Example

The following example establishes an SSH2 session on switch engineering1:

```
# ssh2 admin@engineering1
```

The following example establishes an SSH2 session with the switch named X460-G2 over TCP port 2050 with compression enabled:

```
# ssh2 compression on port 2050 admin@X460-G2
```

History

This command was first available in ExtremeXOS 11.2.

Changes to **cipher**, as well as the addition of **mac**, were first available in ExtremeXOS 15.7.1.

Ciphers and MACs that are unsupported in OpenSSH 8.1p1 were removed in ExtremeXOS 30.7.

Platform Availability

This command is available on all platforms with the SSH2 module installed.

start orchestration mlag

```
start orchestration mlag peer_name
```

Description

Ensures configuration consistency between controlling bridges (CBs) in a redundant CB environment with bridge port extenders (BPEs) or peers.

Syntax Description

orchestration	Specifies entering orchestration mode in redundant controlling bridges environments.
mlag	Multi-switch link aggregation group.
<i>peer_name</i>	Name of the MLAG peer controlling bridge (switch).

Default

N/A.

Usage Guidelines

With redundant CBs attached to each BPE, the associated extended port configuration must be identical on each controlling bridge. To reduce the configuration complexity and to minimize the risk of inconsistency, you can use orchestration mode so that any configuration commands are now checkpointed to the MLAG peer switch. Before entering orchestration mode, ensure that any configuration parameters (connecting ports for the BPEs, VLAN names, numbers etc.) are the same for both CBs.

After entering this orchestration mode, like the existing virtual-router mode, the configuration prompt changes indicating that commands issued are within this context:

```
(orchestration bottom) X670G2-48x-4q.4 #
```



Note

MLAG peer checkpoint status must be 'up' to enter orchestration mode.



Note

Enabling orchestration mode on both MLAG peers at same time is not recommended. If orchestration mode is enabled on both MLAG peers and you execute commands on both switches at same time, the switch may abort execution of the command.

Example

The following example enables orchestration mode on the controlling bridge (switch) with its MLAG peer "bottom":

```
# start orchestration mlag "bottom"
```

History

This command was first available in ExtremeXOS 22.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

start process

```
start process name {msm slot}
```

Description

Starts the specified process on the switch. Used to restart a process after it has been terminated.

Syntax Description

<i>name</i>	Specifies the name of the process to start. You can start the following processes: bgp, eaps, exsshd, isis, lldp, netLogin, netTools, ospf, snmpMaster, snmpSubagent, telnetd, tftpd, vrrp, xml
<i>slot</i>	On a SummitStack, specifies the node's slot number. The number is a value from 1 to 8.

Default

N/A.

Usage Guidelines

Use this command after you have stopped a process and you want to restart it. To stop a process, use the [terminate process](#) command.

You are unable to start a process that is already running. If you try to start a currently running process, an error message similar to the following appears:

```
Error: Process telnetd already exists!
```

Depending on the software version running on your switch and the type of switch you have, you can restart different or additional processes. To see which processes you can restart, enter start process followed by [Tab]. The switch displays a list of available processes.

To display the status of ExtremeXOS processes on the switch, including how many times a process has been restarted, use the [show process {name} {detail} {description} {slotslotid}](#) command.

You can also use the [start process](#) command when upgrading a software modular package. For more information, see the section *Upgrading a Modular Software Package* in the [Switch Engine 32.3 User Guide](#).

SummitStack Only

You can issue this command only from the master node. If you issue this command from any other node, the following message appears:

```
Error: Processes created by user can only be started on the primary node slot.
```



Note

After you stop a process, do not change the configuration on the switch until you start the process again. A new process loads the configuration that was saved prior to stopping the process. Changes made between a process termination and a process start are lost, and error messages can result when you start the new process.

Example

The following restarts the process tftpd:

```
start process tftpd
```

History

This command was first available in ExtremeXOS 11.0.

Support for restarting the Link Layer Discovery Protocol (lldp), Open Shortest Path First (ospf), and network login (netLogin) processes was added in ExtremeXOS 11.3.

Support for restarting the Border Gateway Protocol (bgp) was added in ExtremeXOS 11.4.

Support for restarting netTools was added in ExtremeXOS 12.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

start vm

```
start vm vm_name
```

Description

Starts a virtual machine (VM) process.

Syntax Description

vm	Designates starting a vm process.
<i>vm_name</i>	Specifies the VM name.

Default

N/A.

Usage Guidelines

The effect of this command is not persistent. The initial VM state after boot is determined by the VM's enabled/disabled configuration state.

To stop a VM process, use the command `stop vm vm_name [forceful | graceful]`.

The Extreme Insight feature requires the Solid State Storage Device SSD-120.

Example

The following example starts the VM process "vm1":

```
# start vm vm1
```

History

This command was first available in ExtremeXOS 30.3.

Platform Availability

This command is available on all platforms that support the Extreme Insight for Guest VMs feature and have a Core license installed. For a list of platforms that support the Insight feature and for information about licenses, see the [Switch Engine 32.3 Feature License Requirements](#).

stop orchestration

stop {**orchestration**}

Description

Exits orchestration mode, which is a mode used in a redundant controlling bridges environment with bridge port extenders (BPE) or peers.

Syntax Description

orchestration	Specifies exiting orchestration mode in redundant controlling bridges environments with MLAG peers.
----------------------	---

Default

N/A.

Usage Guidelines

After exiting this mode, the configuration prompt changes indicating that commands issued are no longer within this context. See the following example.

Example

The following example exits orchestration mode:

```
(orchestration mlag peer) X460G2-24t-10G4.2 # stop orchestration
X460G2-24t-10G4.2 #
```

History

This command was first available in ExtremeXOS 22.5.

Platform Availability

This command is available on ExtremeSwitching 5420 and 5520 series switches.

stop vm

stop **vm** *vm_name* [**forceful** | **graceful**]

Description

Stops a virtual machine (VM) process.

Syntax Description

vm	Designates stopping a vm process.
<i>vm_name</i>	Specifies the VM name.
forceful	Forcefully terminates the VM.
graceful	Gracefully shutdowns the VM if possible (default).

Default

By default, the VM is shutdown gracefully, if possible.

Usage Guidelines

The effect of this command is not persistent. The initial VM state after boot is determined by the VM's enabled/disabled configuration state. A VM is forcefully stopped if it does not shutdown gracefully within 60 seconds

To start a VM process, use the command `start vm vm_name`.

The Extreme Integrated Application Hosting (IAH) feature requires the Solid State Storage Device SSD-120.

Example

The following example stops the VM process "vm1" forcefully:

```
# stop vm vm1 forcefully
```

History

This command was first available in ExtremeXOS 30.3.

Platform Availability

This command is available on all platforms that support the Extreme IAH feature and have a Core license installed. For a list of platforms that support the IAH feature and for information about licenses, see the [Switch Engine 32.3 Feature License Requirements](#).

suspend vm

```
suspend vm vm_name
```

Description

Suspends a virtual machine (VM).

Syntax Description

vm	Specifies a VM.
<i>vm_name</i>	Specifies the VM name to suspend.

Default

N/A.

Usage Guidelines

Suspending a VM pauses, or freezes, the CPU state of the guest VM. The VM remains in an active state, and committed resources (such as memory and ports) remain committed. This state is not persistent across switch reboots or the virtMgr process restarts. Such events function as "power off resets" to the guest operating system.

To resume the VM, use the `resume vm vm_name` command.

Example

The following example pauses the VM "vm1":

```
# suspend vm vm1
```

History

This command was first available in ExtremeXOS 30.7.

Platform Availability

This command is available on all platforms that support the Integrated Application Hosting (IAH) feature and have a Core license installed. For a list of platforms that support the IAH feature and for information about licenses, see the [Switch Engine 32.3 Feature License Requirements](#).

synchronize

```
synchronize {slot slotid}
```

Description

The synchronize command replicates all saved images and configurations from the master node to the backup or target node on a SummitStack.

Syntax Description

<i>slotid</i>	Specifies the target node that has to be synchronized with the master node. If the slotid is omitted, the target is the backup node. Note: This parameter is available only with SummitStack.
---------------	---

Default

If you do not specify a node, the backup node is synchroized with the master.

Usage Guidelines



Note

You cannot synchronize a master switch running a version earlier than ExtremeXOS 22.2 to a target node running ExtremeXOS 22.2 or later. The command fails and the following error message appears: `Error: the target slot's partitions are not compatible with the Master's for synchronize.`

This command:

- Reboots the backup node or target node to prepare it for synchronizing with the master node.
- Performs a binary copy of the master node to the backup or target node, including the primary and secondary software images, all configurations and policies, and temporary files.
- Reboots the backup or target node after replication is complete.

During a synchronization, half of the switch fabric is lost. When the master node finishes replicating its configurations and images to the backup or target node, the full switch fabric is restored.

To use this command, make sure your SummitStack is running ExtremeXOS 12.0 or later.

When you install a different backup node, you are not prompted to synchronize the images and the configurations from the master. If not synchronized, the backup node uses its image and the master's configuration. This image/configuration mismatch will likely cause the switch to operate differently after failover. Use the [synchronize](#) command to replicate all saved images and configurations from the master to the backup node.

If you have not saved your runtime configuration, you are prompted to save it when you use this command. A message similar to the following appears:

```
Do you want to save configuration changes to primary.cfg? (y or n)
```

Enter `y` to save the configuration and continue with synchronizing the nodes. Enter `n` to cancel the operation. If you enter `y`, messages similar to the following appear:

```
Saving configuration on primary ..... done! Synchronizing configuration to backup ..  
done!
```

After the configuration has been saved and replicated to the backup node, synchronization begins.

After the initial reboot, if the backup node is not available or does not respond within 120 seconds, the synchronize operation fails.

Use the `show switch {detail}` command to verify that the backup node is in sync with the master node.

While using the command `synchronize {slot slot-number}` if the slot number is provided, that slot is the target of the synchronize operation. If the slot number is not provided, the backup node is synchronized. This command can be executed only on the master node.

The synchronize command preserves the following stacking configurations on the target node:

- slot number
- master-capable configuration
- alternate IP address, subnetwork mask, and default gateway
- priority
- license restriction

Thus a synchronized node comes up in the same place in the active topology that it occupied before the synchronize command was issued.

If the synchronizing switch cannot determine the stacking support configuration on the target switch, the following message is displayed:

```
Error: Information for the target switch is temporarily unavailable. Please retry the command.
```

Example

The following example assumes you have already saved your runtime configuration. Using `synchronize` replicates all saved images and configurations from the master node to the backup node.

After you enter the command, status messages similar to the following appear:

```
# synchronize
Synchronize will reboot the backup, then overwrite all code images
and configs with a copy from the master.
Synchronize currently requires ExtremeXOS version 11 or greater on
the backup
DO NOT interrupt synchronize, the backup may become unbootable!
OK to continue? (y/n) Yes
Rebooting Backup...
NOTE: The command line is locked during synchronize
synchronizing...
synchronizing nvram...
synchronizing nvram...
synchronizing nvram...
synchronizing nvram...
synchronizing nvram...
synchronizing nvram...
synchronizing XOS...
[=====] 100% XOS
Synchronize complete - rebooting backup ...
```


History

This command was first available in ExtremeXOS 11.0.

The slot parameter was added to support SummitStack in ExtremeXOS 12.0.

Platform Availability

This command is available only on SummitStacks.

synchronize stacking

```
synchronize stacking {node-address node_address | slot slot_number}
```

Description

This command copies certain NVRAM based configuration parameters to the target node.

Syntax Description

<i>node_address</i>	Specifies the MAC address of a node in the stack. To view the MAC addresses for all nodes in a stack, enter the show stacking command.
<i>slot_number</i>	Specifies the slot number of the target node. To view the slot numbers, enter the show stacking command.

Default

N/A.

Usage Guidelines

This command synchronizes the following NVRAM-based parameters:

- Stacking mode.
- Stack MAC address.
- Failsafe account and password.
- Failsafe account access point permissions (whether the failsafe account is allowed over the stacking links, console port, or management port).
- The selected partition.

These parameters are copied from the executing node's NVRAM to the target node's NVRAM.

Example

Example for the synchronize stacking command output:

```
Slot-2 Stack.3 > synchronize stacking slot 3  
Are you sure you want to synchronize the specified slot with this slot's stacking
```

```

configuration? (y/n) Yes
Synchronized configuration will take effect at the next reboot of the specified node(s).
Slot-2 Stack.4 >
Slot-2 Stack.4 > synchronize stacking node 00:04:96:27:87:10
Are you sure you want to synchronize the specified node with this node's stacking
configuration? (y/n) Yes
Synchronized configuration will take effect at the next reboot of the specified node(s).
Slot-2 Stack.5 >
Slot-2 Stack.5 > synchronize stacking
Are you sure you want to synchronize all remote nodes with this node's stacking
configuration? (y/n) Yes
Synchronized configuration will take effect at the next reboot of the specified node(s).
Slot-2 Stack.6 >

```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available with all licenses and platforms that support the SummitStack feature. For information about which licenses and platforms support the SummitStack feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

telnet slot

```
telnet slot slot-number {no-auto-login}
```

Description

Allows you to Telnet from any source node to any other target node when both nodes are in the same active topology.

Syntax Description

<i>slot-number</i>	Specifies the number of a slot that is currently occupied by the intended target node.
no-auto-login	Blocks auto-login. You are prompted for login credentials. Without using this option, after you are logged on to one slot (node), you can Telnet to another slot with the same user credentials.

Default

N/A.

Usage Guidelines

After the target node accepts the related TCP connection, when using Telnet to connect to another node, you are not prompted for your user name or password. You are logged in to the same account (with corresponding rights) with which you accessed the originating slot.

If the failsafe account is used, user ID and password authentication takes place on the specified node. Otherwise, authentication takes place on the master node, regardless of the source and target nodes used.

Telnet must be enabled on the SummitStack for this command to function.

Example

The following command accesses the node in slot 2:

```
Slot-1 Stack.3 # telnet slot 2
Entering character mode
Escape character is '^]'.

telnet session telnet0 on /dev/ptyb0

ExtremeXOS
Copyright (C) 1996-2017 Extreme Networks. All rights reserved.
This product is protected by one or more US patents listed at http://www.extremenetworks.com/patents along with their foreign counterparts.
=====

You are connected to a Backup node. Only a limited command set is supported.
You may use "telnet slot <slot_number>" to connect to the Master node to access
the full set of commands.
Press the <tab> or '?' key at any time for completions.
Remember to save your configuration changes.

Slot-2 Stack.1 >
```

History

This command was first available in ExtremeXOS 12.0.

Automatic user authentication from the source node to the target node added in ExtremeXOS 22.3.

The **no-auto-login** option was added in ExtremeXOS 22.5.

Platform Availability

This command is available only on SummitStack.

telnet

```
telnet {vr vr_name} [host_name | remote_ip] {port}
```

Description

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

Syntax Description

vr	Specifies use of a virtual router. Note: User-created VRs are supported only on the platforms listed for this feature in the Switch Engine 32.3 Feature License Requirements document.
<i>vr_name</i>	Specifies the name of the virtual router.
<i>host_name</i>	Specifies the name of the host.
<i>remote_ip</i>	Specifies the IP address of the host.
<i>port</i>	Specifies a TCP port number. The default is port 23.

Default

- Telnet—enabled.
- Virtual router—Uses all virtual routers on the switch for outgoing Telnet requests.
- Port—23.

Usage Guidelines

Only VT100 emulation is supported.

Before you can start an outgoing 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 want to connect to. Check the user manual supplied with the Telnet facility if you are unsure of how to do this. Although the switch accepts IPv6 connections, you can only Telnet from the switch to another device with an IPv4 address.

You must configure DNS in order to use the *host_name* option.

Host Name and Remote IP Address Character Restrictions

This section provides information about the characters supported by the switch for host names and remote IP addresses.

When specifying a host name or remote IP address, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-) Permitted only for host names
- Underscore (_) Permitted only for host names
- Colon (:)

When naming or configuring an IP address for your network server, remember the requirements listed above.

Virtual Router Requirements

The `vr_name` option specifies the name of the virtual router. The valid virtual router names at system boot-up are [VR-Mgmt](#), [VR-Control](#), and [VR-Default](#); however, you can only Telnet on VR-Mgmt and VR-Default. In ExtremeXOS 10.1, the valid virtual routers are VR-0, VR-1, and VR-2 respectively, and Telnet used VR-0 by default. For more information about virtual routers, see [Virtual Routers](#).

Example

The following command starts a Telnet client communication to the host at IP address 123.45.67.8:

```
telnet 123.45.67.8
```

The following command starts a Telnet client communication with a host named sales:

```
telnet sales
```

History

This command was first available in ExtremeXOS 10.1.

Support for the following virtual routers was added in ExtremeXOS 11.0: VR-Mgmt and VR-Default.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

terminate process

```
terminate process name [forceful | graceful]
```

Description

Terminates the specified process on the switch.

Syntax Description

<i>name</i>	Specifies the name of the process to terminate. You can terminate the following processes: bgp, eaps, exsshd, isis, lldp, netLogin, netTools, ntp, ospf, ospfv3, snmpMaster,snmpSubagent, telnetd, thttpd, tftpd, vrrp, and xmld.
forceful	Specifies a forceful termination.
graceful	Specifies a graceful termination.

Default

N/A.

Usage Guidelines

If recommended by Extreme Networks Technical Support personnel, you can stop a running process.

The forceful option quickly terminates a process on demand. Unlike the graceful option, the process is immediately shutdown without any of the normal process cleanup. The status of the operation is displayed on the console. After a successful forceful termination of a process, a message similar to the following appears:

```
Forceful termination success for snmpMaster
```

The graceful option terminates the process by allowing it to close all opened connections, notify peers on the network, and other types of process cleanup. After this phase, the process is finally terminated. After a successful graceful termination of a process, a message similar to the following appears:

```
Successful graceful termination for snmpSubagent
```



Note

In a stack, starting, restarting, or terminating user-created processes in any node other than the primary node is not allowed.

Example

The following initiates a graceful termination of the process tftpd:

```
terminate process tftpd graceful
```

History

This command was first available in ExtremeXOS 11.0.

Support for terminating the Link Layer Discovery Protocol (lldp), network login (netLogin), and Open Shortest Path First (ospf) processes was added in ExtremeXOS 11.3.

Support for terminating the Border Gateway Protocol (bgp) and Ethernet Automatic Protection Switching (eaps) processes was added in ExtremeXOS 11.4.

Support for terminating the MultiProtocol Label Switch (mpls) and Virtual Router Redundancy Protocol (vrrp) processes was added in ExtremeXOS 11.6.

Support for terminating netTools was added in ExtremeXOS 12.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

terminate vpex ztp

```
terminate vpex ztp
```

Description

Prevents automatic configuration of Extended Edge Switching topologies.

Syntax Description

vpex	Specifies Virtual Port Extender (VPEX).
ztp	Zero Touch Provision (ZTP). Automated process for configuring Extended Edge Switching topologies.

Default

If you do not run this command, automatic configuration (ZTP) occurs when an unconfigured controlling bridge (CB) (new, out of the shipping box, or manually unconfigured) is rebooted with attached bridge port extenders (BPEs).

Usage Guidelines

This command allows you to prevent automatic configuration from running. Automatic configuration is re-enabled after rebooting or running the `restart process ztpstack` command.



Note

Alternatively, to prevent automatic configuration from running on Extended Edge Switching topologies, you can perform some configuration on the CB prior to attaching BPEs or have a `default.xsf` (or other similar file) file in place.

To view the automatic configuration (ZTP) status, use the `show vpex` command.

Example

The following example prevents automatic configuration from running:

```
# terminate vpex ztp
```

One of the following messages appears:

```
VPEX ZTP successfully terminated. Allow up to 15 seconds for the status to be reflected
in the "show vpex" output.
VPEX ZTP not running. No need to terminate.
```

History

This command was first available in ExtremeXOS 32.3.

Platform Availability

This command is available on ExtremeSwitching 5420 and 5520 series switches.

tftp

```
tftp [ ip-address | host-name ] { -v vr_name } { -b block_size } [ -g |
  -p ] [ -l local-file { -r remote-file } | -r remote-file { -l local-
  file } ]
```

Description

Allows you to TFTP from the current command line interface session to a TFTP server.

Syntax Description

<i>ip-address</i>	Specifies the IP address of the TFTP server.
<i>host-name</i>	Specifies the name of the remote host.
<i>vr_name</i>	Specifies the name of the virtual router. Note: User-created VRs are supported only on the platforms listed for this feature in the Switch Engine 32.3 Feature License Requirements document.
<i>block_size</i>	Specifies the data block size, excluding TFTP header. Data block size ranges from 24-65,000 bytes. Default is 14,00 bytes.
-g	Gets the specified file from the TFTP server and copies it to the local host.
-p	Puts the specified file from the local host and copies it to the TFTP server.
<i>local-file</i>	Specifies the name of the file (configuration file, policy file) on the local host.
<i>remote-file</i>	Specifies the name of the file on the remote host.

Default

If you do not specify a virtual router, *VR-Mgmt*, and then *VR-Default* (if needed), is used.

If you do not specify otherwise, the switch downloads or uploads the file from the switch local file system (/usr/local/tmp)..

If you do not specify block size, the default value is 14,00 bytes.

Usage Guidelines

NetASCII and mail file type formats are not supported.

TFTP Server Requirements

Extreme Networks recommends using a TFTP server that supports blocksize negotiation (as described in *RFC 2348*, TFTP Blocksize Option) to enable faster file downloads and larger file downloads. If the TFTP server does not support blocksize negotiation, the file size is limited to 32MB. Older TFTP servers that do not support blocksize negotiation have additional implementation limits that may decrease the maximum file size to only 16MB, which may be too small to install ExtremeXOS images.

If your TFTP server does not support blocksize negotiation, the switch displays a message similar to the following when you attempt a get (-g) or put (-p) operation:

Note: The `blocksize` option is not supported by the remote TFTP server. Without this option, the maximum file transfer size is limited to 32MB. Some older TFTP servers may be limited to 16MB file.

Using TFTP

Use TFTP to download a previously saved configuration file or policy file from the TFTP server to the switch. When you download a file, this command does not automatically apply it to the switch. You must specify that the downloaded file be applied to the switch. For example, if you download a configuration file, run the `use configuration` command to apply the saved configuration on the next reboot. You must run the `reboot` command to activate the new configuration. If you download a policy file, run the `refresh policy` command to reprocess the text file and update the policy database.

You also use TFTP to upload a saved configuration file or policy file from the switch to the TFTP server.

If your download from the TFTP server to the switch is successful, the switch displays a message similar to the following:

```
Downloading megtest2.cfg to switch... done!
```

If your upload from the switch to the TFTP server is successful, the switch displays a message similar to the following:

```
Uploading megtest1.cfg to TFTPHost ... done!
```

Up to eight active TFTP sessions can run on the switch concurrently.

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

Host Name and Remote IP Address Character Restrictions

When specifying a host name or remote IP address, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-) Permitted only for host names
- Underscore (_) Permitted only for host names
- Colon (:)

Local and Remote File Name Character Restrictions

When specifying a local or remote file name, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-)
- Underscore (_)
- Slash (/) permitted only for remote files

Virtual Router Requirements

The `vr_name` option specifies the name of the virtual router. The valid virtual router names at system boot-up are VR-Mgmt, [VR-Control](#), and VR-Default; however, you can only TFTP on VR-Mgmt and VR-Default. In ExtremeXOS 10.1, the valid virtual routers are VR-0, VR-1, and VR-2 respectively. For more information about virtual routers, see the *Virtual Routers* chapter in the [Switch Engine 32.3 User Guide](#).

Internal Memory and Core Dump Files

Core dump files have a `.gz` file extension. The file name format is: `core.process-name.pid.gz` where `process-name` indicates the name of the process that failed and `pid` is the numerical identifier of that process.

If you configure and enable the switch to send core dump (debug) information to internal memory, specify the internal-memory file path `/usr/local/tmp` to transfer those files from internal memory to a TFTP server.

If the switch has not saved any debug files, you cannot transfer other files to or from internal memory. For example if you attempt to transfer a configuration file from the switch to the internal memory, the switch displays a message similar to the following:

```
Update: Error: tftp transfer to /usr/local/tmp is not allowed.
```

For information about configuring and sending core dump information, see the `configure debug core-dumps [off | directory_path] save debug tracefiles directory_path` and commands.

For more detailed information about core dump files, see [Troubleshooting](#).

Other Useful Commands

On the ExtremeSwitching series switches and SummitStack, use the [download bootrom](#) command to upgrade the BootROM. This command utilizes TFTP to transfer the BootROM image file from your TFTP server to the switch. Only upgrade the BootROM when asked to do so by an Extreme Networks technical representative. For more information about this command, see [download bootrom](#).

To upgrade the image, run the [download image](#) command. This command uses TFTP to transfer the software image file from your TFTP server to the switch. For more information about this command, see [download image](#).

Example

The following example downloads the configuration file named `XOS1.cfg` from the TFTP server with an IP address of 10.123.45.67:

```
# tftp 10.123.45.67 -v "VR-Default" -g -r XOS1.cfg
```

The following example uploads the configuration file named `XOS2.cfg` to the TFTP server with an IP address of 10.123.45.67:

```
# tftp 10.123.45.67 -v "VR-Default" -p -r XOS2.cfg
```

The following example downloads a policy file to a USB storage device:

```
# tftp 10.1.2.3 -g -l /usr/local/ext/test.pol -r august23.pol
```

History

This command was first available in ExtremeXOS 10.1.

The memory card option was added in ExtremeXOS 11.1.

The internal-memory option was added in ExtremeXOS 11.4.

Support for USB 2.0 storage devices was added in ExtremeXOS 12.5.3.

Path name support added in ExtremeXOS 15.5.1.

Block size support was added in ExtremeXOS 15.7.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

tftp get

```
tftp get [ ip-address | host-name ] { vr vr_name } { block-size  
      block_size } remote-file local-file } {force-overwrite}
```

Description

Allows you to use TFTP from the current command line interface session to copy the file from a TFTP server and copy it to a local host, including the switch, internal memory card, compact flash card, or USB 2.0 storage device.

Syntax Description

<i>host-name</i>	Specifies the name of the remote host.
<i>ip-address</i>	Specifies the IP address of the TFTP server.
<i>vr_name</i>	Specifies the name of the virtual router. NOTE: User-created VRs are supported only on the platforms listed for this feature in the Switch Engine 32.3 Feature License Requirements document.
<i>block_size</i>	Specifies the data block size, excluding TFTP header. Data block size ranges from 24-65000 bytes.
<i>local_file</i>	Specifies the name of the file (configuration file, policy file) on the local host.
force-overwrite	Specifies the switch to automatically overwrite an existing file.

Default

If you do not specify a virtual router, *VR-Mgmt*, and then *VR-Default* (if needed), is used. If you transfer a file with a name that already exists on the system, the switch prompts you to overwrite the existing file.

If you do not specify the internal memory card or a removable storage device, the switch downloads or uploads the file from the switch local file system.

If you do not specify block size, the default value is 1400 bytes.

Usage Guidelines

NetASCII and mail file type formats are not supported.

By default, the switch prompts you to overwrite an existing file. For example, if you have a file named test.cfg on the switch and download a file named test.cfg from a TFTP server, the switch displays a message similar to the following:

```
test.cfg already exists, do you want to overwrite it? (y/n)
```

Enter y to download the file and overwrite the existing file. Enter n to cancel this action.

If you successfully download the file, the switch displays a message similar to the following:

```
Downloading test.cfg to switch... done!
```

If you cancel this action, the switch displays a message similar to the following:

```
Tftp download aborted.
```

If you specify the force-overwrite parameter, the switch automatically overwrites an existing file. For example, if you have a file named test.cfg on the switch and download a file named test.cfg from a TFTP server, the switch automatically overrides the existing file. If you successfully download the file, the switch displays a message similar to the following:

```
Downloading test.cfg to switch... done!
```

This command was introduced to simplify using TFTP to transfer configuration, policy, and if configured, core dump files from the switch to the TFTP server. You can continue to use the original [tftp](#) command introduced in ExtremeXOS 10.1.

For more information about TFTP, including:

- TFTP server requirements.
- How to use TFTP.
- Host name and remote IP address character restrictions.
- Local and remote filename character restrictions.
- Virtual router requirements.
- Internal memory and core dump files.
- Other useful commands.

See the [tftp](#) command.

Example

The following command retrieves and transfers the file test.pol from a TFTP server with an IP address of 10.1.2.3 and renames the file august23.pol when transferred to a removable storage device:

```
tftp get 10.1.2.3 vr "VR-Mgmt" test.pol /usr/local/ext august23.pol
```

The following command retrieves the configuration file named meg-upload.cfg from a TFTP server with an IP address of 10.10.10.10:

```
tftp get 10.10.10.10 vr "VR-Mgmt" meg_upload.cfg
```

History

This command was first available in ExtremeXOS 11.4.

Support for USB 2.0 storage devices was added in ExtremeXOS 12.5.3.

Pathname support for local files was added in ExtremeXOS 15.5.1.

Block size support was added in ExtremeXOS 15.7.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

tftp put

```
tftp put [ ip-address | host-name] {vr vr_name} {block-size  
      block_size}local-file { remote-file}
```

Description

Allows you to use TFTP from the current command line interface session to copy the file from the local host, including the switch, internal memory card, compact flash card, or USB 2.0 storage device and put it on a TFTP server.

Syntax Description

<i>host-name</i>	Specifies the name of the remote host.
<i>ip-address</i>	Specifies the IP address of the TFTP server.
<i>vr_name</i>	Specifies the name of the virtual router. NOTE: User-created VRs are supported only on the platforms listed for this feature in the Switch Engine 32.3 Feature License Requirements document.
<i>block_size</i>	Specifies the data block size, excluding TFTP header. Data block size ranges from 24-65000 bytes.
<i>local-file</i>	Specifies the name of the file (configuration file, policy file) on the local host.
<i>remote-file</i>	Specifies the name of the file on the remote host.

Default

If you do not specify a virtual router, *VR-Mgmt*, and then *VR-Default* (if needed), is used.

If you do not specify a virtual router, VR-Mgmt, and then VR-Default (if needed), is used.

If you do not specify the internal memory card or a removable storage device, the switch downloads or uploads the file from the switch local file system.

If you do not specify block size, the default value is 1400 bytes.

Usage Guidelines

This command is supported by IPv4 only.

NetASCII and mail file type formats are not supported.

This command was introduced to simplify using TFTP to transfer configuration, policy, and if configured, core dump files from the switch to the TFTP server. You can continue to use the original `tftp` command introduced in ExtremeXOS 10.1.

For more information about TFTP, including:

- TFTP server requirements.
- How to use TFTP.
- Host name and remote IP address character restrictions.
- Local and remote filename character restrictions.
- Virtual router requirements.
- Internal memory and core dump files.
- Other useful commands.

See the `tftp` command.

Example

The following command transfers a saved, not currently used configuration file named XOS1.cfg from the switch to the TFTP server:

```
tftp put 10.123.45.67 vr "VR-Mgmt" XOS1.cfg
```

History

This command was first available in ExtremeXOS 11.4.

Support for USB 2.0 storage devices was added in ExtremeXOS 12.5.3.

Block size support was added in ExtremeXOS 15.7.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

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 process, sampled every 5 seconds.

You can change the display by typing a character while the display is active. The following table displays the supported commands.

Table 46: TOP Interactive Command Display Options

Key	Action
P	Sort process list by CPU utilization
T	Sort process list by time usage
N	Sort process list by number (process ID)
M	Sort process list by memory usage
q [Ctrl] + c	Exit the top program

For more detailed information about the top command including display options, command fields, and command usage, please refer to your UNIX documentation.

Example

The following command displays the real-time CPU utilization information by process:

```
top
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

traceroute

```
traceroute {vr vrid} {ipv4 host} {ipv6 host} {ttl number} {from from}
           {[port port] | icmp}
```

Description

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

Syntax Description

vr	Specifies a virtual router. Note: User-created VRs are supported only on the platforms listed for this feature in the Switch Engine 32.3 Feature License Requirements document.
<i>vrid</i>	Specifies which virtual router. Note: User-created VRs are supported only on the platforms listed for this feature in the Switch Engine 32.3 Feature License Requirements document.
ipv4	Specifies IPv4 transport.
ipv6	Specifies IPv6 transport.
<i>host</i>	Specifies the host of the destination endstation.
ttl <i>number</i>	Configures the switch to trace up to the time-to-live number of the switch.
from <i>from</i>	Uses the specified source address in the <i>ICMP</i> packet. If not specified, the address of the transmitting interface is used.
port <i>port</i>	Specifies the UDP port number.
icmp	Configures the switch to send ICMP echo messages to trace the routed path between the switch and a destination endstation.

Default

N/A.

Usage Guidelines

Use this command to trace the routed path between the switch and a destination endstation.

Each router along the path is displayed.

Example

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

```
traceroute 123.45.67.8
```

The following is sample output that displays when the traceroute fails:

```
traceroute to 10.209.10.37, 30 hops max
1  0.0.0.0                                * !u          * !u          * !u
--- Packet Response/Error Flags ---
(*) No response, (!N) ICMP network unreachable, (!H) ICMP host unreachable,
(!P) ICMP protocol unreachable, (!F) ICMP fragmentation needed,
(!S) ICMP source route failed, (!u) Transmit error, network unreachable,
(!f) Transmit error, fragmentation needed, (!t) General transmit error
```

History

This command was first available in ExtremeXOS 10.1.

The IPv6 variable was added in ExtremeXOS 11.2.

The display when the command fails was added in ExtremeXOS 11.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

traceroute mac port

```
traceroute mac mac {up-end-point} port port {domain} domain_name
                {association} association_name {ttl ttl}
```

Description

Allows you to send out a Link Trace Message (LTM) for the specified MA from the MEP configured on the port for the specified MAC address to the end of the MA.

Syntax Description

<i>mac</i>	Enter the unique system MAC address on the port configured as a MEP for the specified MA. Enter this value in the format XX:XX:XX:XX:XX:XX.
up-end-point	Use this keyword to force the LTM to be send from an UP MEP if both a DOWN MEP and an UP MEP are configured on the same port.
<i>port</i>	Enter the port number of the MEP from which you are issuing the LTM.

domain	Enter this keyword.
<i>domain_name</i>	IEEE 802.1ag or ITU-T Y.1731 association name.
association	Enter this keyword.
<i>association_name</i>	Enter the name of the association from which you are issuing the ping.
ttl	Enter this keyword.
<i>ttl</i>	Enter the upper limit of MIPs the LTM can pass prior to reaching its destination.

Default

TTL default value is 64.

Usage Guidelines

Use this command to send an LTM from the MEP on the port for the given MAC address. If no MEP is configured on the port, the system returns an error message.

If both an UP and DOWN MEP are configured on the same port, the system uses the DOWN MEP. If you want to use the UP MEP in this situation, enter the up-end-point keyword. After you issue the command, the system prints out the route the LTM message took.

Each MIP along the route passes the LTM along only in the direction of the path and sends a packet back to the originating MAC notifying that it passed the LTM. If the destination MAC type is configured as unicast on the association to which this MEP belongs to, link trace replies will not be received from any of the MIPs configured on the intermediate switches. If there is a MIP on the switch that originated the trace route, the MIP sends a link trace reply.

Example

The following commands send an LTM:

1. A trace route invoked from a customer device CE1 to another customer device CE3 connected through an *MPLS* cloud (MTU1 -' PE1 'PE3), where a VPLS MIP is configured to encode a system-name, will have a response as follows:

```
(debug) Switch # traceroute mac 00:04:96:28:02:15 port 1 "extr_cfm5" "extr_ma"
Send out Link Trace Message(LTM), collecting responses [press Ctrl-C to abort].
TTL  CFM Source MAC      Reply  Reply Mac              Port ID
=====
63   00:04:96:1e:6d:40    I F-f-  00:04:96:1e:6d:40    o--    1:8
62   00:04:96:1e:6d:40    E F-f-  00:04:96:1e:6d:40    o--    vp100:MTU-1
61   00:04:96:1e:16:10    I F-f-  00:04:96:1e:16:10    o--    vp100:PE-1
60   00:04:96:1e:16:10    E F-f-  00:04:96:1e:16:10    o--    vp100:PE-1
59   00:04:96:1e:14:90    I F-f-  00:04:96:1e:14:90    o--    vp100:PE-3
58   00:04:96:1e:14:90    E F-f-  00:04:96:1e:14:90    o--    1:8
57   00:04:96:28:02:15    I -h--  00:04:96:28:02:15    o--    1
=====
Reply Flags: (I) Ingress, (E) Egress, (F) FwdYes, (h) RlyHit, (f) RlyFDB
Flags: (o) Ok, (d) Down, (b) Blocked
```

2. A trace route Invoked within an MPLS Cloud from MTU1 to PE3 (MTU1 -' PE1 'PE3), where a VPLS MIP is configured to encode a private-ip, will have a response as follows:

```
(debug) Switch # traceroute mac 00:04:96:1e:14:90 port 1:8 extr_cfm2 "extr_ma"
Send out Link Trace Message(LTM), collecting responses [press Ctrl-C to abort].
TTL  CFM Source MAC      Reply  Reply Mac      Port ID
=====
63   00:04:96:1e:6d:40    E F-f-  00:04:96:1e:6d:40  o--  vp100:3.3.3.3
62   00:04:96:1e:16:10    I F-f-  00:04:96:1e:16:10  o--  vp100:1.1.1.1
61   00:04:96:1e:16:10    E F-f-  00:04:96:1e:16:10  o--  vp100:5.5.5.5
60   00:04:96:1e:14:90    I F-f-  00:04:96:1e:14:90  o--  vp100:3.3.3.3
59   00:04:96:1e:14:90    E -h--  00:04:96:1e:14:90  o--  1:8
=====
Reply Flags: (I) Ingress, (E) Egress, (F) FwdYes, (h) RlyHit, (f) RlyFDB
Flags: (o) Ok, (d) Down, (b) Blocked
```

If in PE1 alone, a VPLS MIP is configured to encode a system name, the response will be as follows:

```
(debug) Switch # traceroute mac 00:04:96:1e:14:90 port 1:8 extr_cfm2 "extr_ma"
Send out Link Trace Message(LTM), collecting responses [press Ctrl-C to abort].
TTL  CFM Source MAC      Reply  Reply Mac      Port ID
=====
63   00:04:96:1e:6d:40    E F-f-  00:04:96:1e:6d:40  o--  vp100:3.3.3.3
62   00:04:96:1e:16:10    I F-f-  00:04:96:1e:16:10  o--  vp100:PE1
61   00:04:96:1e:16:10    E F-f-  00:04:96:1e:16:10  o--  vp100:PE1
60   00:04:96:1e:14:90    I F-f-  00:04:96:1e:14:90  o--  vp100:3.3.3.3
59   00:04:96:1e:14:90    E -h--  00:04:96:1e:14:90  o--  1:8
=====
Reply Flags: (I) Ingress, (E) Egress, (F) FwdYes, (h) RlyHit, (f) RlyFDB
Flags: (o) Ok, (d) Down, (b) Blocked
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

traceroute mpls lsp

```
traceroute mpls lsp [lsp_name | any host | prefix ipNetmask] {reply-mode
    [ip | ip-router-alert]} {{from from} {ttl t1} {next-hop hopaddress}}
```

Description

Traces the path an LSP takes for the specified FEC.

Syntax Description

<i>lsp_name</i>	Specifies the LSP on which to send the <i>MPLS</i> echo request.
any	Allows the echo request to be sent over any available LSP.

<i>host</i>	Specifies the FEC using an ipaddress or hostname.
prefix	Specifies a prefix.
<i>ipNetmask</i>	Specifies the prefix address.
reply-mode	Specifies the reply mode for the MPLS echo response.
ip	Requests an IP UDP reply packet. This is the default mode.
ip-router-alert	Requests an IP UDP reply packet with the IP Router Alert option.
<i>from</i>	Specifies the IP address to be used as the source address in the MPLS echo request.
<i>ttl</i>	Specifies the starting TTL hop value. The range is from 1 - 30. The default is 1.
<i>hopaddress</i>	Specifies the next-hop address.

Default

The maximum time-to-live value is 30 seconds.

Usage Guidelines

This command traces the path an LSP takes for the specified FEC. The `traceroute` command, with the **mpls** keyword option, works by repeatedly sending an MPLS echo request (or “MPLS Ping”). The TTL value is incremented for each successive MPLS echo request sent. The sending LSR waits 5 seconds before sending the next MPLS echo request. This operation continues until either the egress LSR for the FEC is reached, the maximum TTL value is reached, or the operation is interrupted. For each response received, the following information is displayed on the console:

- IP address of the replying LSR
- Return code
- Indication of an MPLS echo reply timeout if no response was received

The FEC can be specified using the ipaddress or hostname via the *host* parameter. If the optional **next-hop** is specified, the MPLS echo request is sent along the LSP that traverses the specified node. This option is useful for tracing a specific LSP when multiple LSPs exist to the specified FEC. The *lsp* keyword may be used to specify a named LSP to trace. The selected LSP is specified by the *lsp_name* parameter. The **any** keyword indicates that the switch can trace any available LSP to the specified host.

The optional **reply-mode** keyword is used to specify the reply mode for the MPLS echo response. When the **ip** option is specified, the MPLS echo reply is routed back to the sender in a normal IPv4 packet. When the **ip-router-alert** option is specified, the MPLS echo reply is routed back to the sender in an IPv4 packet with the Router Alert IP option set. Additionally, if the **ip-router-alert** option is specified and the reply route is via an LSP, the Router Alert Label is pushed onto the top of the label stack. If the reply-mode is not specified, the **reply-mode ip** option applies.

The optional **ttl** keyword specifies the starting TTL value in the MPLS echo request packet. Within each router along the path, the TTL value is decremented. When the TTL value reaches zero, the LSR drops the packet and replies with a TTL-expired *ICMP* message. The originating LSR responds by displaying the hop for which the TTL expired. To discover all hops to a destination, the originating router repeats the MPLS echo request and increments the TTL start value by one each time until

the destination is reached. The maximum TTL is 30, so the `traceroute` command terminates if the destination is not reached in 30 hops.

If the `t1` keyword is omitted, the starting TTL value is 1. If you specify a larger starting TTL value, initial hops are excluded from the `traceroute` display. For example, if you specify a start TTL value of 5, the TTL value does not decrement to 0 at the first four routers, so the fifth hop router is the first to appear in the `traceroute` command display.

The `from` keyword is used to specify the source IP address used in the MPLS echo request. This is the IP address used by the target LSR to send the MPLS echo reply. If not specified, the [OSPF](#) router ID is used.

Example

The following example shows a sample display for the `traceroute` command:

```
# traceroute mpls lsp prefix 11.100.100.10/32
traceroute to 11.100.100.10, 30 hops max
 1  11.100.100.8                5 ms      5 ms      2 ms
 2  11.100.100.10              2 ms      1 ms      2 ms
# traceroute mpls lsp lsp598
traceroute to lsp598, 30 hops max
 1  11.100.100.5                6 ms      1 ms      5 ms
 2  11.100.100.9                3 ms      2 ms      2 ms
 3  11.100.100.8                3 ms      4 ms      3 ms
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

unalias

```
unalias alias_name -a
```

Description

Deletes aliases for ExtremeXOS commands.

Syntax Description

<i>alias_name</i>	Specifies the alias name to be deleted.
-a	Deletes all aliases.

Default

N/A

Usage Guidelines

To see a list of your current aliases, use the command [alias](#) on page 87 with no arguments.

When using this command, you can auto-complete an alias name by typing part of the name, and then pressing **[Tab]**.

Example

The following example deletes all aliases:

```
unalias -a
```

History

This command was first available in ExtremeXOS 22.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure access-list

```
unconfigure access-list policy_name {any | ports port_list | vlan
vlan_name} {ingress | egress}
```

Description

Removes a policy file ACL from the specified interface.

Syntax Description

<i>policy_name</i>	Specifies the ACL policy name. The name can be from 1-32 characters long.
<i>port_list</i>	Specifies the ingress or egress port list on which the ACL is applied.
<i>vlan_name</i>	Specifies the <u>VLAN</u> on which the ACL is applied.
ingress	Remove the ACL for packets entering the switch on this interface.
egress	Remove the ACL for packets leaving the switch from this interface

Default

The default direction is ingress.

Usage Guidelines

This command removes ACLs that are contained in ACL policy files. To remove dynamic ACLs, use the following command:

```
configure access-list delete ruleName [ any | vlanvlan_name |  
portsport_list | all] {ingress | egress}
```

To remove all non-dynamic ACLs from all interfaces, do not specify any ports or VLANs.

Example

The following command removes the ACL from port 1:2:

```
unconfigure access-list ports 1:2
```

The following command removes the ACLs from ports 1:2-6:3 and 7:1:

```
unconfigure access-list ports 1:2-6:3,7:1
```

The following command removes the wildcard ACL:

```
unconfigure access-list any
```

The following command removes all ACLs from all the interfaces, including the wildcard ACL:

```
unconfigure access-list
```

History

This command was first available in ExtremeXOS 10.1.

The VLAN option was first available in ExtremeXOS 11.0.

The egress option was first available in ExtremeXOS 11.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

The egress options are available on all platforms.

unconfigure avb

```
unconfigure avb
```

Description

This command is a macro command that can be used to unconfigure all AVB protocols globally on the switch. It is equivalent to issuing the following four commands:

```
unconfigure mvrp  
unconfigure msrp  
unconfigure network-clock gptp  
unconfigure mrp ports all
```

Syntax Description

avb	Audio Video Bridging
------------	----------------------

Default

N/A.

Usage Guidelines

Example

Use this command to unconfigure all AVB protocols globally on the switch.

```
unconfigure avb
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure banner

```
unconfigure banner { after-login | before-login }
```

Description

Unconfigures a specified banner from CLI screens.

Syntax Description

after-login	Specifies the banner that is displayed after login.
before-login	Specifies the banner that is displayed before login.

Default

N/A.

Usage Guidelines

Use this command to unconfigure one of two different types of banners:

- CLI session before login.
- CLI session after login.

If no optional parameters are specified, all configured banners are erased. To delete a specific banner, the before-login or after-login keyword must be used.

Banners can also be cleared by configuring a banner with only a <ret> or \n character.

Example

The following command clears the after-login banner, Welcome to the switch:

```
unconfigure banner after-login [Return]
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure bfd vlan

unconfigure bfd vlan *vlan_name*

Description

Unconfigures BFD settings from a specified VLAN.

Syntax Description

<i>vlan_name</i>	Specifies the VLAN name.
------------------	--------------------------

Default

N/A.

Usage Guidelines

Use this command to unconfigure BFD settings from a specified VLAN.

Example

The following command unconfigures the BFD settings on the VLAN named vlan1:

```
# unconfigure bfd vlan vlan1
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure bootprelay dhcp-agent information check

```
unconfigure bootprelay dhcp-agent information check
```

Description

Disables DHCP relay agent option (option 82) checking.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

In some instances, a DHCP server may not properly handle a DHCP request packet containing a relay agent option. Use this command to disable the switch from preventing DHCP reply packets with invalid or missing relay agent options from being forwarded to the client.

To enable this check, use the following command:

```
configure bootprelay dhcp-agent information check
```

Example

The following example disables the DHCP relay agent option check:

```
unconfigure bootprelay dhcp-agent information check
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure bootprelay dhcp-agent information circuit-id port-information

```
unconfigure bootprelay dhcp-agent information circuit-id port-
information ports [port_list | all]
```

Description

Configures the circuit ID sub-option that identifies the specified ports to use the default value.

Syntax Description

<i>port_list</i>	Specifies a list of one or more ports that are to be configured to use the default value.
all	Specifies that all ports are to be configured to use the default value.

Default

The port_info is encoded as ((slot_number * 1000) + port_number). For example, if the *DHCP* request is received on port 3:12, the default circuit ID port_info value is 3012. On non-slot-based switches, the default circuit ID port_info value is simply the port number.

Usage Guidelines

None.

Example

The following example configures port 1:3 to use the default circuit ID port information value:

```
unconfigure bootprelay dhcp-agent information circuit-id port-information ports 1:3
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure bootprelay dhcp-agent information circuit-id vlan-information

```
unconfigure bootprelay dhcp-agent information circuit-id vlan-
information {vlan} [vlan_name|all]
```

Description

Configures the circuit ID sub-option that identifies the specified VLANs to use the default value.

Syntax Description

<i>vlan_name</i>	Names a VLAN to be configured to use the default value.
all	Specifies that all VLANs are to be configured to use the default value.

Default

N/A.

Usage Guidelines

None.

Example

The following example configures VLAN "blue" to use the default VLAN information for the circuit ID sub-option:

```
unconfigure bootprelay dhcp-agent information circuit-id vlan-information blue
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure bootprelay dhcp-agent information option

```
unconfigure bootprelay dhcp-agent information option
```

Description

Disables the DHCP relay agent option (option 82).

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

To enable the DHCP relay agent option (option 82), use the following command:

```
configure bootrelay dhcp-agent information option
```

Example

The following example disables the DHCP relay agent option:

```
unconfigure bootrelay dhcp-agent information option
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure bootrelay dhcp-agent information policy

```
unconfigure bootrelay dhcp-agent information policy
```

Description

Unconfigures the DHCP relay agent option (option 82) policy.

Syntax Description

This command has no arguments or variables.

Default

Replace.

Usage Guidelines

Use this command to unconfigure the policy for the relay agent.

Example

The following example unconfigures the DHCP relay agent option 82 policy:

```
unconfigure bootrelay dhcp-agent information policy
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure bootprelay dhcp-agent information remote-id

```
unconfigure bootprelay dhcp-agent information remote-id {vr vrid}
```

Description

Configures the remote ID sub-option to the default value.

Syntax Description

<i>vrid</i>	Specifies the VR on which to configure the remote ID sub-option to the default value.
-------------	---

Default

The switch MAC address.

Usage Guidelines

None.

Example

The following example configures the remote ID sub-option to use the default value on the current VR:

```
configure bootprelay dhcp-agent information remote-id
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure bootprelay dhcp-agent source-vlan

```
unconfigure bootprelay dhcp-agent source-vlan {vr vrid}
```

Description

Unconfigures the source VLAN to use as the source IP address in the BOOTPrelay packet.

Syntax Description

bootprelay	Specifies BOOTPrelay agent information option.
dhcp-agent	Specifies DHCP agent.
source-vlan	Specifies unconfiguring the source loopback VLAN to use as the source IP address in the giaddr field when BOOTPrelay is used for anycast VLAN.
vr	Specifies unconfiguring the VLAN on the specific virtual router ID.
<i>vrid</i>	Specifies the virtual router ID.

Default

If you do not specify a VR, the unconfiguring occurs on VR-Default.

Usage Guidelines

To view the selected source VLAN, use the command `show bootprelay configuration {ipv4 | ipv6} {{vlan vlan_name } | {vr vr_name}}`.

Example

The following example unconfigures the VLAN to use as the source IP address in the BOOTPrelay packet on vr "vr1":

```
# unconfigure bootprelay dhcp-agent source-vlan vr vr1
```

History

This command was first available in ExtremeXOS 30.7.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure bootprelay include-secondary

```
unconfigure bootprelay {ipv4 | ipv6} {vlan} vlan_name include-secondary
```

Description

Removes the specified smart relay configuration that was specified at the VLAN level.

Syntax Description

ipv4	Specifies unconfiguring the DHCPv4 BOOTP Relay service (default).
ipv6	Specifies unconfiguring the DHCPv6 BOOTP Relay service.

vlan	Unconfigures BOOTP relay for a specified VLAN.
<i>vlan_name</i>	Specifies the VLAN name.
include-secondary	Removes the include-secondary configuration for the specified VLAN.

Default

IPv4 is the default relay service.

Usage Guidelines

Use this command to remove the specified smart relay configuration that was specified at the VLAN level.

Example

The following command removes DHCPv4 BOOTP Relay service, and removes the **include-secondary** configuration for the VLAN "vlan 100":

```
unconfigure bootprelay ipv4 vlan vlan_100 include-secondary
```

History

This command was first available in ExtremeXOS 15.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure cfm domain association end-point transmit-interval

```
unconfigure cfm domain domain_name association association_name {ports
port_list end-point [up | down]} transmit-interval
```

Description

Unconfigures the CCM interval of the association or MEP to the default interval.

Syntax Description

<i>domain_name</i>	Specifies the domain associated with the MA.
<i>association_name</i>	IEEE 802.1ag or ITU-T Y.1731 association name.
<i>port_list</i>	Specifies the ports to unconfigure.
up	Enter this variable if you are changing the time interval for sending a CCM on an UP MEP.
down	Enter this variable if you are changing the time interval for sending a CCM on a DOWN MEP.

Default

1000 ms.

Usage Guidelines

Use this command to revert the CCM interval of either the association or the MEP back to the default CCM interval.

Example

The following command changes the interval the UP MEP (previously configured on port 2:4) uses to send CCM messages on the 350 association in the finance domain to the default of 1000 ms:

```
unconfigure cfm domain finance association 350 ports 2:4 end-point up transmit-interval
```

History

This command was first available in ExtremeXOS 12.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure cos-index

```
unconfigure cos-index cos_index [{qosprofile} {ingress-meter} {replace-tos}]
```

Description

This command removes the cos-index from the configuration when no attribute is entered or if the last one is unconfigured.

Syntax Description

<i>cos_index</i>	Class of Service index value.
qosprofile	QoS profile.
ingress-meter	Ingress rate-limiter meter.
replace-tos	Replace TOS value.

Default

N/A.

Usage Guidelines

Use this command to remove the cos-index from the configuration.

Example

```
unconfigure cos-index 50
unconfigure cos-index 51 qosprofile ingress-meter replace-dot1p
unconfigure cos-index 52 qosprofile
unconfigure cos-index 52 ingress-meter
unconfigure cos-index 52 replace-dot1p
```

History

This command was first available in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure diffserv examination

```
unconfigure diffserv examination
```

Description

Disables DiffServ traffic groups.

Syntax Description

N/A.

Default

Disabled.

Usage Guidelines

Use this command to disable DiffServ code point examination.

Example

The following command disables DiffServ code point examination:

```
unconfigure diffserv examination
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure diffserv replacement

```
unconfigure diffserv replacement
```

Description

Resets all DiffServ replacement mappings to the default values.

Syntax Description

N/A.

Default

N/A.

Usage Guidelines

Use this command to reset all DiffServ replacement mappings to default values.

Example

The following command resets the DiffServ replacement mappings to their default values:

```
unconfigure diffserv examination
```

History

This command was first available in ExtremeXOS 11.0.

The ports keyword was first available in ExtremeXOS 12.2.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure eaps port

```
unconfigure eaps eapsDomain [primary | secondary] port
```

Description

Sets the specified port's internal configuration state to INVALID.

Syntax Description

<i>eapsDomain</i>	Specifies the name of an <i>EAPS</i> 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.

To prevent loops in the network, the switch displays by default a warning message and prompts you to unconfigure the specified EAPS primary or secondary ring port. When prompted, do one of the following:

- Enter `y` to unconfigure the specified port.
- Enter `n` or press [Return] to cancel this action.

If you have considerable knowledge and experience with EAPS, you might find the EAPS loop protection warning messages unnecessary. For more information, see the `configure eaps config-warnings off`.

Example

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

```
unconfigureeapseaps_1primary port
```

The switch displays the following warning message and prompts you to confirm this action:

```
WARNING: Unconfiguring the Primary port from the EAPS domain could cause  
a loop in the network! Are you sure you want to unconfigure the Primary  
EAPS Port? (y/n)
```

Enter `y` to continue and unconfigure the EAPS primary ring port. Enter `n` to cancel this action.

The switch displays a similar warning message if you unconfigure the secondary EAPS port.

History

This command was first available in ExtremeXOS 11.0.

The interactive messages were added in ExtremeXOS 11.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure eaps shared-port link-id

```
unconfigure eaps shared-port ports link-id
```

Description

Unconfigures an *EAPS* link ID on a shared port on the switch.

Syntax Description

<i>ports</i>	Specifies the port number of the Common Link port.
--------------	--

Default

N/A.

Usage Guidelines

None.

Example

The following command unconfigures the link ID on shared port 1:1.

```
unconfigure eaps shared-port 1:1 link-id
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms with the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and what licenses are appropriate for this feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

unconfigure eaps shared-port mode

```
unconfigure eaps shared-port ports mode
```

Description

Unconfigures the *EAPS* shared port mode.

Syntax Description

<i>ports</i>	Specifies the port number of the Common Link port.
--------------	--

Default

N/A.

Usage Guidelines

None.

Example

The following command unconfigures the shared port mode on port 1:1:

```
unconfigure eaps shared-port 1:1 mode
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on all platforms with the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and what licenses are appropriate for this feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

unconfigure elrp-client

```
unconfigure elrp-client {vlan_name | dynamic-vlans}
```

Description

Disables a pending one-shot or periodic ELRP request for the specified VLAN or for all dynamic VLANs.

Syntax Description

<i>vlan_name</i>	Specifies a VLAN name.
dynamic-vlans	Removes all ELRP configuration options for dynamically created VLANs.

Default

N/A.

Usage Guidelines

This command disables a pending one-shot or periodic ELRP request for the specified VLAN or dynamic VLANs.

To start one-time, non-periodic ELRP packet transmission on specified ports of a VLAN using a particular count and interval, use one of the following commands:

- [configure elrp-client one-shot](#) on page 447
- [run elrp](#) on page 2425

To configure periodic transmission of ELRP packets, use the [configure elrp-client periodic](#) command.

The ELRP client must be enabled globally in order for it to work on any VLANs. Use the [enable elrp-client](#) command to globally enable the ELRP client.

The ELRP client can be disabled globally so that none of the ELRP VLAN configurations take effect. Use the [disable elrp-client](#) command to globally disable the ELRP client.

Example

The following command disables a pending ELRP request on VLAN elrp1:

```
unconfigure elrp-client elrp1
```

History

This command was first available in ExtremeXOS 11.1.

The **dynamic-vlans** option was added in ExtremeXOS 22.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure elrp-client disable ports

```
unconfigure elrp-client disable-ports
```

Description

Deletes an ELRP exclude port list.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to remove an ELRP exclude port list.

Example

The following example removes the existing ELRP exclude port list:

```
unconfigure elrp-client disable-port
```

History

This command was first available in ExtremeXOS 12.5.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure erps cfm

unconfigure {**erps**} *ring-name* **cfm**

Description

Unconfigure the CFM maintenance association for the ERPS ring.

Syntax Description

<i>ring-name</i>	Alphanumeric string that identifies the ERPS ring.
------------------	--

Default

N/A.

Usage Guidelines

Use this command to unconfigure connectivity fault management (CFM) for the ERPS ring.

Example

The following command unconfigures connectivity fault management on an ERPS ring named “ring1”:

```
unconfigure erps ring1 cfm
```


History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

unconfigure erps neighbor-port

```
unconfigure erps ring-name neighbor-port
```

Description

Delete the ring protection link (RPL) neighbor configuration for the ERPS ring.

Syntax Description

<i>ring-name</i>	Alphanumeric string that identifies the ERPS ring.
------------------	--

Default

N/A.

Usage Guidelines

See Description.

Example

The following command deletes RPL neighbor configuration for the ERPS ring named “ring1”:

```
unconfigure erps ring1 neighbor-port
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

unconfigure erps notify-topology-change

```
unconfigure {erps} ring-name notify-topology-change {eaps} domain_name
```

Description

Delete an *ERPS* sub-ring from the *EAPS* domain.

Syntax Description

<i>ring-name</i>	Alphanumeric string identifying the ERPS sub-ring.
<i>domain_name</i>	Alphanumeric string identifying the EAPS domain.

Default

N/A.

Usage Guidelines

Use this command to delete an ERPS sub-ring from the EAPS domain.

Example

Example output not yet available and will be provided in a future release.

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

unconfigure erps protection-port

```
unconfigure erps ring-name protection-port
```

Description

Delete ring protection link (RPL) owner configuration for the *ERPS* ring.

Syntax Description

<i>ring-name</i>	Alphanumeric string that identifies the ERPS ring.
------------------	--

Default

N/A.

Usage Guidelines

Use this command to delete ring protection link (RPL) owner configuration for the ERPS ring.

Example

The following command deletes RPL owner configuration on an ERPS ring named “ring1”:

```
unconfigure erps ring1 protection-port
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

unconfigure erps ring-ports west

```
unconfigure erps ring-name ring-ports west
```

Description

Delete ring ports on the ERPS ring.

Syntax Description

<i>ring-name</i>	Alphanumeric string that identifies the ERPS ring.
west	Delete the ring port on the west port of the switch.

Default

N/A.

Usage Guidelines

Use this command to delete ring ports on the ERPS ring. Ring ports are the ports of the switch that connect it to the ERPS ring. This command deletes the ring port on the west port of the switch.



Note

On unconfiguring the west port, the node is treated as an interconnected node.

Example

The following command deletes the ring ports on the west port of the switch for an ERPS ring named “ring1”:

```
unconfigure erps ring1 ring-ports west
```

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on all platforms supported in 12.6 and forward that are running ExtremeXOS.

unconfigure icmp

```
unconfigure 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 example resets all ICMP settings to the default values.

```
unconfigure icmp
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure igmp

```
unconfigure 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:

```
unconfigure igmp
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IPv4 multicast feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

unconfigure identity-management list-precedence

```
unconfigure identity-management list-precedence
```

Description

This command allows you to restore the order of list-precedence to factory defaults. The default list-precedence is: greylist blacklist whitelist.

Syntax Description

This command has no arguments or variables.

Default

greylist, blacklist, whitelist.

Usage Guidelines

Use this command to restore the order of list-precedence to factory defaults. The default list-precedence is: greylist blacklist whitelist.

Example

Example output not yet available and will be provided in a future release.

History

This command was first available in ExtremeXOS 15.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure identity-management

```
unconfigure identity-management {[[database memory-size] | [stale-entry aging-time] | [ports] | [kerberos snooping {aging time}]]}
```

Description

Sets the specified identity management configuration parameter to the default values.

Syntax Description

database memory-size	Sets the identity management database size to the default value.
stale-entry aging-time	Sets the stale-entry aging-time to the default value.
ports	Removes all ports from the identity management port list.
kerberos snooping aging time	Sets the kerberos snooping aging time to the default value (none).

Default

N/A.

Usage Guidelines

If no configuration parameters are specified, all configuration parameters are set to the default values.

Example

The following command sets all identity management configuration parameters to the default values:

```
* Switch.4 # unconfigure identity-management
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure igmp snooping vlan ports set join-limit

```
unconfigure igmp snooping {vlan} vlan_name ports port_list set join-  
limit
```

Description

Removes the join limit set on VLAN ports.

Syntax Description

<i>vlan_name</i>	Specifies a VLAN name.
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

No limit.

Usage Guidelines

None.

Example

The following command removes the join limit for port 2:1 in the Default VLAN:

```
unconfigure igmp snooping "Default" ports 2:1 set join-limit
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the *IGMP* snooping feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

unconfigure igmp ssm-map

```
unconfigure igmp ssm-map {{vr} vr-name}
```

Description

Unconfigures all SSM mappings on the virtual router.

Syntax Description

<i>vr-name</i>	Specifies a virtual router name. If the VR name is omitted, the switch uses the VR specified by the current CLI VR context.
----------------	---

Default

N/A.

Usage Guidelines

None.

Example

The following command deletes all *IGMP*-SSM mappings on the virtual router "xyz":

```
unconfigure igmp ssm-map vr xyz
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the IPv4 multicast feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

unconfigure inline-power classification

```
unconfigure inline-power classification ports port_list
```

Description

This command configures *PoE* port-level classification power-up mode for Extreme Networks PoE devices that do not support 802.3bt when connecting to switches with 60W/90W PoE ports. ExtremeSwitching platforms support per-port basis configuration.

Syntax Description

classification	Classification power-up mode.
ports	Specifies the port.
<i>port_list</i>	Port list separated by a comma or - .

Default

Depends on the maximum classification level supported by the platform. For example, the 5720 Series has a default classification level of 802.3bt-type4, while 5320 Series has a default classification level of 802.3at.

Usage Guidelines

This command restores port-level classification power-up mode to default based on port type.

Example

```
# show inline-power configuration ports 20-26
Port  Config   Operator Limit  Priority  Detection           Classification
Label
20     Enabled   60000 mW      Low      802.3af-only        802.3pre-bt
N
21     Enabled   60000 mW      Low      802.3af-only        802.3bt-type3
N
22     Enabled   60000 mW      Low      802.3af-only        802.3bt-type3
N
23     Enabled   45000 mW      Low      2-point legacy-and-802.3af 802.3bt-type3
N
24     Enabled   60000 mW      Low      802.3af-only        802.3at
N
x435-port10
25     Enabled   90000 mW      Low      4-point legacy-and-802.3af 802.3pre-bt
N
26     Enabled   90000 mW      Low      802.3af-only        802.3bt-type3
N
```

History

This command was first available in ExtremeXOS 31.3.

Platform Availability

This command is available on:

PoE+

- ExtremeSwitching 5420F-8W-16P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420F-24P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420F-16MW-32P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420F-16W-32P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420F-48P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420F-48P-4XL—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420M-16MW-32P-4YE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5320-48P-8XE—ExtremeXOS 31.6 and later.
- ExtremeSwitching 5320-24P-8XE—ExtremeXOS 31.6 and later.
- ExtremeSwitching 5320-16P-4XE—ExtremeXOS 31.7 and later.
- ExtremeSwitching 5320-16P-4XE-DC—ExtremeXOS 31.7 and later.

PoE++

- ExtremeSwitching 5520-24W—ExtremeXOS 31.1 and later.
- ExtremeSwitching 5520-48W—ExtremeXOS 31.1 and later.
- ExtremeSwitching 5520-12MW-36W—ExtremeXOS 31.1 and later.
- ExtremeSwitching 5420F-8W-16P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420F-16MW-32P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420F-16W-32P-4XE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420M-24W-4YE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420M-16MW-32P-4YE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5420M-48W-4YE—ExtremeXOS 31.3 and later.
- ExtremeSwitching 5720-24MW—Switch Engine 32.1 and later.
- ExtremeSwitching 5720-24MXW—Switch Engine 32.1 and later.
- ExtremeSwitching 5720-48MW—Switch Engine 32.1 and later.
- ExtremeSwitching 5720-48MXW—Switch Engine 32.1 and later.

unconfigure inline-power detection ports

```
unconfigure inline-power detection ports [port_list | all]
```

Description

This command unconfigures *PoE* device detection mode for ExtremeSwitching series switches and SummitStacks. ExtremeSwitching platforms support per-port basis configuration.

Syntax Description

<i>port_list</i>	Port list separated by a comma or - .
all	All ports.

Default

Default is **802.3af-only** detection.

Usage Guidelines

None.

History

This command was first available in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure inline-power disconnect-precedence

```
unconfigure inline-power disconnect-precedence
```

Description

Unconfigures the disconnect precedence setting and returns the switch to the default disconnect precedence value of deny port.

Syntax Description

This command has no arguments or variables.

Default

Deny-port.

Usage Guidelines

You configure this parameter for the entire switch; you cannot configure this per slot or per port.

Unconfigures the *PoE* disconnect precedence previously set for the switch and returns the disconnect precedence to the default value of deny port. Deny port denies power to the next PD that requests inline power from the slot when the inline power budget for the switch or slot is reached, regardless of the inline power port priority.

Example

The following command resets the switch to the PoE disconnect precedence value, which is deny port:

```
unconfigure inline-power disconnect-precedence
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on the PoE devices listed in [Extreme Networks PoE Devices](#).

unconfigure inline-power operator-limit ports

```
unconfigure inline-power operator-limit [class-based] ports [all |  
    port_list]
```

Description

Unconfigures the *PoE* operator limit setting and resets the power limit allowed for PDs connected to the specified ports to the default values.

Syntax Description

operator-limit	Specifies to use the default operator limit based on the port's class type (for example, 30W, 60W, 90W).
class-based	Specifies to use the configured operator-limit value for the port(s) instead of the powered device's class.
all	Specifies all ports.
<i>port_list</i>	Specifies one or more slots and ports.

Default

PoE—15400 mW.

PoE+—30000 mW.

Usage Guidelines

This command unconfigures any previously configured operator limit for the specified ports. It resets the maximum power that any PD can draw to 15400 mW for PoE and 30000 mW for PoE+.

Example

The following command resets the limit on ports 3 to 6 on a switch to the default value of 15400 mW:

```
unconfigure inline-power operator-limit ports 3-6
```



The following command configures **class-based** on port 1:

```
configure inline-power operator-limit class-based ports 1
```

History

This command was first available in ExtremeXOS 11.1.

PoE+ was added in ExtremeXOS 12.5.



The **class-based** option was added in ExtremeXOS 32.3.

Platform Availability

This command is available on the PoE devices listed in [Extreme Networks PoE Devices](#).

unconfigure inline-power priority ports

```
unconfigure inline-power priority ports [all | port_list]
```

Description

Unconfigures the *PoE* priority on the specified ports, and returns the ports to the default PoE port priority value of low.

Syntax Description

all	Specifies all ports.
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

Low.

Usage Guidelines

Use this to reset the PoE port priority on specified ports on switches to the default value of low.

If there are multiple ports on the switch at the same priority level (either configured, or by default) and one of the ports must have power withdrawn because of excessive power demands, those ports with the lower port number are powered first. The higher port numbers have power withdrawn first in the case of equal PoE port priorities.

Example

The following command resets the PoE priority on ports 4 – 6 to low:

```
unconfigure inline-power priority ports 4-6
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on the PoE devices listed in [Extreme Networks PoE Devices](#).

unconfigure inline-power usage-threshold

```
unconfigure inline-power usage-threshold
```

Description

Unconfigures the inline power usage alarm threshold and returns threshold to the default value of 70%.

Syntax Description

This command has no arguments or variables.

Default

70.

Usage Guidelines

This command unconfigures the *PoE* usage threshold setting for initiating *SNMP* event and EMS messages and returns the switch's inline power usage threshold for to 70%. The system initiates an event and message once that percentage of the budgeted power is being used.

On stand-alone switches, this PoE threshold applies to the entire switch.

The system generates an additional SNMP event and EMS message once the power usage falls below the threshold again; once the condition clears.

Example

The following command resets the inline power usage alarm threshold to 70%:

```
unconfigure inline-power usage-threshold
```

History

This command was first available in ExtremeXOS 11.1

Platform Availability

This command is available on the PoE devices listed in [Extreme Networks PoE Devices](#).

unconfigure iparp

```
unconfigure iparp
```

Description

Resets the following to their default values:

- IP ARP timeout
- Maximum ARP entries
- Maximum ARP pending entries
- ARP checking
- ARP refresh

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following example resets IP ARP timeout to its default value:

```
unconfigure iparp
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure ip-fix

```
unconfigure ip-fix
```

Description

Unconfigures IPFIX globally.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to unconfigure IPFIX globally by removing all port and collector configuration and disabling IPFIX on all ports.

Example

The following command removes all IPFIX configuration:

```
unconfigure ip-fix
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure ip-fix flow-key

```
unconfigure ip-fix flow-key
```

Description

Unconfigures IPFIX flow key configuration and resets to the default.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to remove the flow-key settings that were configured using the following commands and reset to all the available keys:

```
configure ip-fix flow-key ipv4
```

```
configure ip-fix flow-key ipv6
```

```
configure ip-fix flow-key nonip
```

Example

The following command removes IPFIX flow key settings:

```
unconfigure ip-fix flow-key
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure ip-fix ip-address

```
unconfigure ip-fix ip-address
```

Description

Unconfigures the collector settings.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to unconfigure the collector settings and reset to the default.

Example

The following command returns to the default:

```
unconfigure ip-fix ip-address
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure ip-fix ports

```
unconfigure ip-fix ports port_list
```

Description

Unconfigures IPFIX on a port or group of ports.

Syntax Description

<i>port_list</i>	Specifies the ports.
------------------	----------------------

Default

N/A.

Usage Guidelines

Use this command to unconfigure IPFIX on a port or group of ports. This restores the configuration to the defaults for those ports. The global enable/disable of IPFIX is not affected by this command.

Example

The following command unconfigures IPFIX on port 2:

```
unconfigure ip-fix ports 2
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure ip-fix ports flow-key mask

```
unconfigure ip-fix ports port_list flow-key mask
```

Description

Unconfigures the IPv4 and IPv6 masks.

Syntax Description

<i>port_list</i>	Specifies the ports.
------------------	----------------------

Default

N/A.

Usage Guidelines

Use this command to remove masks for the IPv4 and IPv6 source and destination address fields on ports. These masks were defined using one or more of the following commands:

```
configure ip-fix ports flow-key ipv4 mask ipaddress
```

```
configure ip-fix ports flow-key ipv6 mask ipaddress
```

Example

The following command removes a mask on port 2:1:

```
unconfigure ip-fix ports 2:1 flow-key mask
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure ip-fix source ip-address

```
unconfigure ip-fix source ip-address
```

Description

Unconfigures the source IP address used to communicate to the collector.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to unconfigure the source ipaddress and reset to the default.

Example

The following command returns to the default:

```
unconfigure ip-fix source ip-address
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure iproute priority

```
unconfigure iproute {ipv4} priority [all | blackhole | bootp | ebgp
| ibgp | icmp | isis | isis-level-1 | isis-level-1-external | isis-
level-2 | isis-level-2-external | mpls | ospf-as-external | ospf-
extern1 | ospf-extern2 | ospf-inter | ospf-intra | rip | static |
evpn | ospf-inter-vr] {vr vrname}
```

Description

Unconfigures the priority for all IP routes from one or all route origin types.

Syntax Description

all	Specifies all route origins.
blackhole	Specifies the blackhole route.
bootp	Specifies BOOTP.

ebgp	Specifies E- <i>BGP</i> routes.
ibgp	Specifies I-BGP routes.
icmp	Specifies <i>ICMP</i> .
isis	Specifies IS-IS and applies only to blackhole routes installed for summary addresses.
isis-level-1	Specifies IS-IS Level 1 routing.
isis-level-1-external	Specifies IS-IS Level 1 External routing.
isis-level-2	Specifies IS-IS Level 2 routing.
isis-level-2-external	Specifies IS-IS Level 2 External routing.
mpls	Specifies <i>MPLS</i> routing.
ospf-as-external	Specifies <i>OSPF</i> as External 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 <i>RIP</i> .
static	Specifies static routes.
evpn	Specifies EVPN routes.
ospf-inter-vr	Specifies OSPF inter-VR routes.
<i>vrname</i>	Specifies a VR or VRF name.

Default

N/A.

Usage Guidelines

Unconfiguring the IP route priorities sets the priority back to its default priority.

Default Route Priorities

The following table lists the default priorities that apply after you run this command.

Table 47: Relative Route Priorities

Route Origin	Priority
Direct	10
MPLS	20
Blackhole	50
Static	1100
HostMobility	1150

Table 47: Relative Route Priorities (continued)

Route Origin	Priority
ICMP	1200
EVPN	1698
Autopeering	1699
EBGP	1700
IBGP	1900
OSPFIntra	2200
OSPFInter	2300
IS-IS	2350
IS-IS L1	2360
IS-IS L2	2370
RIP	2400
OSPFAsExt	3100
OSPF External 1	3200
OSPF External 2	3300
IS-IS L1 Ext	3400
IS-IS L2 Ext	3500
BOOTP	5000
OSPFInterVR	5010

Example

The following example returns the IP route priority for all route origins to the default values:

```
# unconfigure iproute priority all
```

History

This command was first available in ExtremeXOS 12.1.2.

The **evpn** option was added in ExtremeXOS 30.7.

The **ospf-inter-vr** option was added in ExtremeXOS 32.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure iproute ipv6 priority

```
unconfigure iproute ipv6 priority [all | blackhole | icmp | isis |
isis-level-1 | isis-level-1-external | isis-level-2 | isis-level-2-
external | ospfv3-as-external | ospfv3-extern1 | ospfv3-extern2 |
ospfv3-inter | ospfv3-intra | ripng | static] {vr vr_name}
```

Description

Resets the priority for all IPv6 routes from one or all route origin types to the default values.

Syntax Description

all	Specifies all route origins.
blackhole	Specifies the blackhole route.
icmp	Specifies <i>ICMP</i> .
isis	Specifies IS-IS and applies only to blackhole routes installed for summary addresses.
isis-level-1	Specifies IS-IS Level 1 routing.
isis-level-1-external	Specifies IS-IS Level 1 External routing.
isis-level-2	Specifies IS-IS Level 2 routing.
isis-level-2-external	Specifies IS-IS Level 2 External routing.
ospfv3-as-external	Specifies <i>OSPF</i> as External routing.
ospfv3-extern1	Specifies OSPF External 1 routing.
ospfv3-extern2	Specifies OSPF External 2 routing.
ospfv3-inter	Specifies OSPF Inter routing.
ospfv3-intra	Specifies OSPF Intra routing.
ripng	Specifies <i>RIP</i> .
static	Specifies static routes.
<i>vr_name</i>	Specifies a VR or VRF name.

Default

N/A.

Usage Guidelines

The following table lists the default values that apply after you enter this command.

Default Route Priorities

Route Origin	Priority
Direct	10
BlackHole	50
Static	1100
ICMP	1200
OSPF3Intra	2200
OSPF3Inter	2300
IS-IS L1	2360
IS-IS L2	2370
RIPg	2400
<u>OSPFv3</u> ASExt	3100
OSPFv3 Extern1	3200
OSPFv3 Extern2	3300
IS-IS L1 Ext	3400
IS-IS L2 Ext	3500

Example

The following example returns the IPv6 route priority for all route origins to the default values:

```
unconfigure iproute ipv6 priority all
```

History

This command was first available in ExtremeXOS 12.1.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

unconfigure ip-security dhcp-bindings storage filename

```
unconfigure ip-security dhcp-bindings storage filename
```

Description

Disassociates a storage filename for DHCP binding information.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command disassociates a DHCP bindings storage filename that was created on the external server.

History

This command was first available in ExtremeXOS 31.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure ip-security dhcp-snooping information check

```
unconfigure ip-security dhcp-snooping information check
```

Description

Disables the *DHCP* relay agent option (option 82) checking in the server-originated packets.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command disables the checking of the server-originated packets for the presence of option 82 so the packets will be forwarded normally.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure ip-security dhcp-snooping information circuit-id port-information ports

```
unconfigure ip-security dhcp-snooping information circuit-id port-  
information ports [port_list | all]
```

Description

Unconfigures the port information portion of the circuit ID.

Syntax Description

<i>port_list</i>	Specifies the port(s) for which port information of the circuit-ID is unconfigured.
all	Specifies all ports.

Default

The default is all.

Usage Guidelines

This command unconfigures the port information portion of the circuit ID string for the indicated ports thereby restoring it to the default (ifIndex value).

History

This command was first available in ExtremeXOS 12.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure ip-security dhcp-snooping information circuit-id vlan-information

```
unconfigure ip-security dhcp-snooping information circuit-id vlan-  
information [dynamic | {vlan} vlan_name | all]
```

Description

Unconfigures the VLAN info portion of the circuit ID of a VLAN.

Syntax Description

<i>vlan_name</i>	Specifies the VLAN for which VLAN information of the circuit-ID is unconfigured.
all	Specifies all VLANs.
dynamic	Configuration options for dynamically created VLANs.

Default

The default is all.

Usage Guidelines

This command unconfigures the VLAN information portion of the circuit ID of a VLAN, restoring it to the default.

History

This command was first available in ExtremeXOS 12.1.

Dynamic VLAN option added in ExtremeXOS 30.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure ip-security dhcp-snooping information option

```
unconfigure ip-security dhcp-snooping information option
```

Description

Disables the *DHCP* relay agent option (option 82).

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command disables the DHCP relay agent option (option 82), which is inserted into client-originated DHCP packets before they are forwarded to the server.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure ip-security dhcp-snooping information policy

```
unconfigure ip-security dhcp-snooping information policy
```

Description

Unconfigures the *DHCP* relay agent option (option 82) policy.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command unconfigures the DHCP relay agent option information policy to the default value of replace.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure ip-security dhcp-snooping information remote-id

```
unconfigure ip-security dhcp-snooping information remote-id
```

Description

Removes the *DHCP* relay agent remote ID.

Syntax Description

This command has no arguments or variables.

Default

After using this command, or if the remote ID has not been configured, the default remote ID is the switch's MAC address.

Usage Guidelines

This command removes the remote ID set by the `configure ip-security dhcp-snooping information remote-id [system-name | remote-id_info]` command. After running this command, or if the remote ID has never been configured, the default remote ID is the switch's MAC address. However, this default (MAC address) name does not appear in the `show ip-security dhcp-snooping information remote-id` command.

Example

The following command removes the DHCP remote ID:

```
# unconfigure ip-security dhcp-snooping information remote-id
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

Object Missing

This object is not available in the repository.

unconfigure isis area

```
unconfigure isis area area_name {level [1|2]}
```

Description

This command resets most configurable parameters for the specified router process to the default values.

Syntax Description

<i>area_name</i>	Specifies the router process to be unconfigured.
level [1 2]	Specifies either level 1 or level 2.

Default

N/A.

Usage Guidelines

This command does not delete interfaces from the router process, but it does disable them. The system ID and IS type are not changed. Where appropriate, the default values apply to level 1, level 2, and both IPv4 and IPv6.

Example

The following command resets the area configuration parameters for areax:

```
unconfigure isis area areax
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Premier license as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

unconfigure isis vlan

```
unconfigure isis [vlan all | {vlan} vlan_name] {level [1|2]}
```

Description

This command resets all configurable interface parameters to the defaults on one or all [VLANs](#).

Syntax Description

vlan all	Unconfigures IS-IS for all VLANs.
<i>vlan_name</i>	Specifies a single VLAN for which IS-IS is unconfigured.
level [1 2]	Specifies either level 1 or level 2.

Default

N/A.

Usage Guidelines

None.

Example

The following command resets IS-IS configuration parameters for SJvlan:

```
unconfigure isis SJvlan
```

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on platforms with a Premier license as described in the [*Switch Engine 32.3 Feature License Requirements*](#) document.

unconfigure l2vpn dot1q ethertype

```
unconfigure l2vpn [vpws vpws_name | vpls vpls_name] dot1q ethertype
```

Description

Resets the ethertype setting for the specified VPLS or VPWS.

Syntax Description

<i>vpls_name</i>	Identifies the VPLS within the switch (character string).
<i>vpws_name</i>	Identifies the VPWS within the switch (character string).

Default

N/A.

Usage Guidelines

The setting is changed back to the value displayed in the `show dot1q` command.

The `l2vpn` keyword is introduced in ExtremeXOS Release 12.4 and is required when resetting the ethertype for a VPWS. For backward compatibility, the **l2vpn** keyword is optional when resetting the ethertype for a VPLS. However, this keyword will be required in a future release, so we recommend that you use this keyword for new configurations and scripts.

Example

The following command changes the ethertype setting for the specified VPLS to the value displayed in the `show dot1q` command:

```
unconfigure l2vpn vpls my_vpls dot1q ethertype
```

History

This command was first available in ExtremeXOS 11.6.

The **l2vpn** and **vpws** keywords were first available in ExtremeXOS 12.4.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

unconfigure l2vpn vpls redundancy

```
unconfigure {l2vpn} vpls vpls_name redundancy [eaps | esrp | stp]
```

Description

Disassociates the VPLS instance from [EAPS](#), an [ESRP](#) domain, or [STP](#).

Syntax Description

<i>vpls_name</i>	Specifies the VPLS for which you are removing protection.
eaps	Disassociates the VPLS instance from EAPS.
esrp	Disassociates the VPLS instance from the ESRP domain.
stp	Disassociates the VPLS instance from STP.

Default

Redundancy disabled.

Usage Guidelines

The **l2vpn** keyword is introduced in ExtremeXOS Release 12.4. For backward compatibility, the **l2vpn** keyword is optional. However, this keyword will be required in a future release, so Extreme Networks recommends that you use this keyword for new configurations and scripts.

Example

The following command disassociates the VPLS instance from ESRP:

```
unconfigure l2vpn vpls vpls1 redundancy esrp
```

History

This command was first available in ExtremeXOS 12.1.

The **l2vpn** keyword and the **STP** option were added in ExtremeXOS 12.4.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

unconfigure ldap domains

```
unconfigure ldap domains
```

Description

This command deletes all LDAP domains, and thereby all LDAP servers and other LDAP configurations done for those domains.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to delete all LDAP related configuration from the switch.

Example

The following command deletes all LDAP configurations, LDAP servers and LDAP domains.

```
Switch.25 # unconfigure ldap domains
```

History

This command was first available in ExtremeXOS 15.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure lldp

```
unconfigure lldp {ports [all | port_list]}
```

Description

Leaves [LLDP](#) enabled and configured; restores the LLDP timer default values.

Syntax Description

all	Specifies all ports on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

N/A.

Usage Guidelines

When you issue the global **unconfigure lldp**, only the LLDP timers are reset to default values. All the configured TLVs remain on the ports remain, and LLDP remains enabled.

When you use the keyword **ports**, the TLVs for each port are returned to the five default TLVs. LLDP remains enabled.

Example

The following command restores LLDP factory default TLVs for ports 1:4 to 1:8:

```
unconfigure lldp ports 1:4 - 1:8
```

History

This command was first available in ExtremeXOS 11.2.

The keyword **port** was changed to **ports** in ExtremeXOS 11.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure log filter

```
unconfigure log filter filter_name
```

Description

Resets the log filter to its default values; removes all filter items.

Syntax Description

<i>filter_name</i>	Specifies the log filter to unconfigure.
--------------------	--

Default

N/A.

Usage Guidelines

If the filter name specified is DefaultFilter, this command restores the configuration of DefaultFilter back to its original settings.

If the filter name specified is not DefaultFilter, this command sets the filter to have no events configured and therefore, no incidents will pass. This is the configuration of a newly created filter that was not copied from an existing one.

See the [delete log filter](#) command for information about deleting a filter.

Example

The following command sets the log filter myFilter to stop passing any events:

```
unconfigure log filter myFilter
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available all platforms.

unconfigure log target format

```
unconfigure log target [console | memory-buffer | nvram | session |  
  syslog [all | ipaddress {udp-port {udp_port}} | ipPort | ipaddress  
  tls-port {tls_port} ] {vr vr_name} {local0...local17}] format
```

Description

Resets the log target format to its default values.

Syntax Description

console	Specifies the console display format.
memory-buffer	Specifies the switch memory buffer format.
nvr am	Specifies the switch NVRAM format.
session	Specifies the current session (including console display) format.
syslog	Specifies a syslog target format.
all	Specifies all remote syslog servers.

<i>ipaddress</i>	Specifies the syslog IP address.
<i>ipPort</i>	Specifies the UDP port number for the syslog target.
tls_port	Specifies remote Syslog server Transport Layer Security (TLS) for connection type.
<i>tls_port</i>	TLS port number.
<i>vr_name</i>	Specifies the virtual router that can reach the server IP address. Note: User-created VRs are supported only on the platforms listed for this feature in the Switch Engine 32.3 Feature License Requirements document.
local0 ... local7	Specifies the local syslog facility.
format	Specifies that the format for the target will be reset to the default value.

Default

When a target format is unconfigured, it is reset to the default values.

The following defaults apply to console display, memory buffer, NVRAM, and session targets:

- timestamp—hundredths
- date—mm-dd-yyyy
- severity—on
- event-name—condition
- host-name—off
- sequence-number—off
- process-name—off
- process-id—off
- source-line—off

The following defaults apply to syslog targets (per RFC 3164):

- timestamp—seconds
- date—mmm-dd
- severity—on
- event-name—none
- host-name—off
- sequence-number—off
- process-name—off
- process-id—off
- source-line—off

Usage Guidelines

Use this command to reset the target format to the default format.

Example

The following command sets the log format for the target session (the current session) to the default:

```
unconfigure log target session format
```

History

This command was first available in ExtremeXOS 10.1.

The **udp-port** parameter was added in ExtremeXOS 21.1.

Transport Layer Security (TLS) option added in ExtremeXOS 22.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure meter

```
unconfigure meter metername ports [port_group | port_list ]
```

Description

This command allows the unconfiguring a per-port meter for one set of ports but not another.

Syntax Description

ports	Meter configuration is applicable to ports in the specified <i>port_group</i> or <i>port_list</i> .
<i>port_group</i>	Port group name.
<i>port_list</i>	Port list separated by a comma.

Default

N/A.

Usage Guidelines

With global meters, there is no need for an unconfigure command as the user can simply delete the global meter to remove the configuration. However, the per-port meters are created by default (e.g. *ingmeter0*, *ingmeter1*...) and cannot be deleted. This command allows the unconfiguring a per-port meter for one set of ports but not another.

Example

```
unconfigure meter ingmeter1 ports ingMtrGrpA
```

History

This command was first available in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure mlag peer interval

```
unconfigure mlag peer peer_name interval
```

Description

Unconfigures the length of time between health check hello packets.

Syntax Description

<i>peer_name</i>	Specifies an alpha numeric string identifying the peer.
------------------	---

Default

The interval default is 1000 milliseconds.

Usage Guidelines

Use this command to unconfigure the length of time between health check hello packets exchanged between MLAG peer switches and reset to the default.

Example

The following command unconfigures the interval on the switch101 peer. switch:

```
# unconfigure mlag peer switch101 interval
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure mlag peer ipaddress

```
unconfigure mlag peer peer_name ipaddress
```

Description

Unconfigures an peer switch IP address from an MLAG structure.

Syntax Description

<i>peer_name</i>	Specifies an alpha numeric string identifying the MLAG peer.
------------------	--

Default

N/A.

Usage Guidelines

Use this command to disassociate an MLAG peer structure with an MLAG peer switch IP address.

Example

The following command disassociates the MLAG peer structure switch101 with the MLAG peer switch IP address:

```
# unconfigure mlag peer switch101 ipaddress
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure mld

```
unconfigure mld
```

Description

Resets all MLD settings to their default values and clears the MLD group table.

Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command resets all MLD settings to their default values and clears the MLD group table:

```
unconfigure mld
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 multicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

unconfigure mld ssm-map

```
unconfigure mld ssm-map {{vr}vr_name }
```

Description

Deletes all MLD SSM Mapping entries on a VR.

Syntax Description

vr vr_name	Specifies the virtual router name.
-------------------	------------------------------------

Default

Disabled.

Usage Guidelines

Use this command to delete all MLD SSM Mapping entries on a specified VR.

Example

The following command deletes MLD SSM mapping on VR1 :

```
unconfigure mld ssm-map vr vr1
```

History

This command was first available in ExtremeXOS 15.5.

Platform Availability

This command is available on the platforms listed for the IPv6 multicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

unconfigure mpls exp examination

```
unconfigure mpls exp examination [all | {{value} value}]
```

Description

Resets the [QoS](#) profile assigned to the EXP value back to the default QoS profile.

Syntax Description

<i>value</i>	Specifies the EXP value whose QoS profile is reset.
--------------	---

Default

The QoS profile matches the EXP value + 1.

Usage Guidelines

This command resets the QoS profile assigned to the EXP value (defined by the value keyword and argument) back to the default QoS profile. If the **all** option is specified, all EXP values are reset back to their default QoS profiles. By default, the QoS profile matches the EXP value + 1. That is, EXP value of 0 is mapped to QoS profile qp1, EXP value of 1 is mapped to QoS profile qp2, etc. This configuration has switch-wide significance.

Example

Use the following command to restore the QoS profile of EXP value 5 to its default setting:

```
unconfigure mpls exp examination value 5
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

unconfigure mpls exp replacement

```
unconfigure mpls exp replacement [all | {{qosprofile} qosprofile}]
```

Description

Resets the EXP value assigned to the [QoS](#) profile back to the default EXP value.

Syntax Description

<i>qosprofile</i>	Specifies the QoS profile whose EXP replacement value is unconfigured.
-------------------	--

Default

The EXP value matches the QoS profile -1.

Usage Guidelines

This command resets the EXP value assigned to the QoS defined by *qosprofile* back to the default EXP value. If the **all** option is specified, all QoS profiles are reset back to their default EXP values. By default, the EXP value matches the QoS profile - 1. That is, QoS profile qp1 is mapped to EXP value of 0, QoS profile qp2 is mapped to EXP value of 1, etc. This configuration has switch-wide significance.

Example

Use the following command to restore all EXP values to their default setting:

```
unconfigure mpls exp replacement all
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

unconfigure mpls vlan

```
unconfigure mpls {{vlan} vlan_name}
```

Description

Resets [MPLS](#) configuration parameters to the default settings. This command does not delete the [VLAN](#) from MPLS.

Syntax Description

<i>vlan_name</i>	Specifies the VLANs for which MPLS is unconfigured.
------------------	---

Default

N/A.

Usage Guidelines

This command resets all MPLS configuration parameters for the specified VLAN to their default values. It does not delete the VLAN from MPLS. These parameters include the enable state for LDP and RSVP-TE, the bandwidth reserved for RSVP-TE LSPs, RSVP-TE timers, and the RSVP-TE metric. MPLS does not have to be disabled to unconfigure a specific VLAN.

Example

The following command resets MPLS configuration parameters to the default settings for a single VLAN:

```
unconfigure mpls vlan boone
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

unconfigure mpls

```
unconfigure 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 deletes all VLANs from MPLS and resets all MPLS configuration parameters to their default values. The parameters that are reset include the LSR ID, all LDP-specific settings, all RSVP-TE-

specific settings, all RSVP-TE reserved bandwidth, and all EXP Qos Profile mappings. MPLS must be disabled to unconfigure MPLS.



Note

MPLS must be disabled to globally unconfigure MPLS.

Example

The following command resets MPLS configuration parameters to the default settings:

```
unconfigure mpls
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support MPLS as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

unconfigure mrp ports timers

```
unconfigure mrp ports [port_list | all] {timers {refresh} {join} {leave}  
  {leave-all} {periodic}}
```

Description

Unconfigure MRP timers, or only reset the MRP timer values to default if the **timer** keyword is specified.

Syntax Description

mrp	Multiple Registration Protocol.
ports	Ports on which MRP timers are to be configured.
all	All ports.
timers	Multiple Registration Protocol timers.
refresh	Timer value to use in place of regular leave timer, only in cases when leave-all is received or sent.
join	The time interval to delay sending MRP advertisements.
leave	The time interval to wait in the leaving state before transitioning to the empty state.
leave-all	The time interval used to control the frequency of "leave all" messages.
periodic	The time interval between two periodic events.

Default

The default values for join, leave, leave-all, are 200, 600, and 10000, respectively. The default values for join, leave, leave-all, periodic and extended-refresh timers are 200, 600, 10000, 1000, and 0 milliseconds, respectively.

Usage Guidelines

Use this command to unconfigure MRP timers, or only reset the MRP timer values to default if the **timer** keyword is specified. If none of the timers are specified, this command resets all three timers to the default values. The default values for the join, leave, and leave-all timers are 200, 600, and 10000 ms respectively.

Example

```
unconfigure mrp ports all
unconfigure mrp ports all timers
unconfigure mrp ports all timers join
```

History

This command was first available in ExtremeXOS 15.3.

The **extended-refresh** and **periodic timer** options were added in 15.3.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure msdp sa-cache-server

```
unconfigure msdp sa-cache-server {vr vrname}
```

Description

Removes the MSDP SA cache server.

Syntax Description

<i>vrname</i>	Specifies the name of the virtual router on which the MSDP cache server is configured. If a virtual router name is not specified, it is extracted from the current CLI context.
---------------	---

Default

By default, the router does not send SA request messages to its MSDP peers when a new member joins a group and wants to receive multicast traffic. The new member simply waits to receive SA messages, which eventually arrive.

Usage Guidelines

Use this command to remove the MSDP SA cache server you specified with the `configure msdp sa-cache-server` command.

Example

The following command removes the MSDP SA cache server:

```
unconfigure msdp sa-cache-server
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the MSDP feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

unconfigure msrp

```
unconfigure msrp {ports [port_list | all]}
```

Description

Disables MSRP and removes all configuration. If a list of ports is specified, MSRP is disabled and the related configuration is removed only on the ports and the system-wide MSRP configuration stays intact.

Syntax Description

msrp	Multiple Stream Registration Protocol.
<i>port_list</i>	List of ports in the switch.
all	All the ports in the switch.

Default

N/A.

Usage Guidelines

Use this command disable MSRP and remove all configuration. If a list of ports is specified, MSRP is disabled and the related configuration is removed only on the ports and the system-wide MSRP configuration stays intact.

Example

```
unconfigure msrp
unconfigure msrp ports all
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on all platforms that support the AVB feature and that have an AVB feature pack license installed. To see which platforms support AVB and for information about obtaining a license, see the [Switch Engine 32.3 Feature License Requirements](#).

unconfigure mstp region

```
unconfigure mstp region
```

Description

Unconfigures the *MSTP* region on the switch and returns all MSTP settings to their default values.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Before you unconfigure an MSTP region, we recommend that you disable all active STPDs in the region. This includes the CIST and any active MSTIs.

After you issue this command, all of the MSTP settings return to their default values, as described below:

- Region Name—This indicates the name of the MSTP region. In the Extreme Networks implementation, the maximum length of the name is 32 characters and can be a combination of alphanumeric characters and underscores (_).
- Format Selector—This indicates a number to identify the format of MSTP BPDUs. The default is 0.
- Revision Level—This identifier is reserved for future use; however, the switch uses and displays a default of 3.

Example

The following command unconfigures the MSTP region on the switch:

```
unconfigure mstp region
```

History

This command was first available in ExtremeXOS 11.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure mvrp stpd

```
unconfigure mvrp stpd
```

Description

Resets the MVRP STP domain.

Syntax Description

mvrp	Multiple <u>VLAN</u> Registration Protocol.
stpd	The STP domain the VLAN is to be associated with. All ports of the domain will be advertised when this VLAN is registered.

Default

s0.

Usage Guidelines

Use this command to reset the STP domain associated with a particular VLAN or all VLANs to default. If a VLAN is specified, the specific VLAN will be associated to the default STP, which is configured using the `configure mvrp stpd stpd_name default` command. If VLAN is not specified, all VLANs are associated to STP domain s0.

Example

The following example illustrates the `unconfigure mvrp stpd` command:

```
unconfigure mvrp stpd
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure mvrp tag

```
unconfigure mvrp tag vlan_tag
```

Description

Resets all MVRP settings for the given *VLAN* id. The *STP* domain, the registrar state machine settings, applicant state machine settings for the given VLAN are reset to default values.

Syntax Description

mvrp	Multiple VLAN Registration Protocol.
<i>vlan_tag</i>	The 802.1Q VLAN ID.

Default

N/A.

Usage Guidelines

Use this command to reset all MVRP settings for the given VLAN id. The STP domain, the registrar state machine settings, and the applicant state machine settings for the given VLAN are reset to default values. All dynamically added ports of the VLAN are removed. If the VLAN was created dynamically, it is removed. If VLAN is not specified, MVRP settings for all VLANs are reset and the dynamic VLAN creation feature is reset to “enabled”.

Example

The following example shows unconfiguring an MVRP:

```
unconfigure mvrp tag 100
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure mvrp

```
unconfigure mvrp
```

Description

Unconfigures MVRP on a switch and all MVRP port and bridge settings.

Syntax Description

mvrp	Multiple <u>VLAN</u> Registration Protocol.
-------------	---

Default

N/A.

Usage Guidelines

Use this command to unconfigure MVRP on a switch. This command unconfigures all MVRP port and bridge settings.

Example

The following command unconfigures MVRP:

```
unconfigure mvrp
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure neighbor-discovery cache

```
unconfigure neighbor-discovery cache {vr vr_name}
```

Description

Resets the neighbor-discovery cache configuration parameters to their default values.

Syntax Description

<i>vr_name</i>	Specifies a VR or VRF.
----------------	------------------------

Default

IPv6 neighbor timeout: 20 minutes

Maximum IPv6 neighbor entries: 1024

Maximum IPv6 neighbor pending entries: 1024

IPv6 neighbor refresh: Enabled

Usage Guidelines

None.

Example

The following example resets the neighbor-discovery cache configuration:

```
unconfigure neighbor-discovery cache
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

unconfigure netlogin

```
unconfigure netlogin [idle-timeout | session-timeout] {dot1x | mac |  
                    web-based | convergence-endpoint}
```

Description

Use this command to reset the maximum number of consecutive seconds an authenticated session may be idle before termination of the session to the default value of 300 seconds.

Syntax Description

idle-timeout	Clear timeout for an idle session.
session-timeout	Clear timeout for an active session.
dot1x	IEEE 802.1X Port-based network access control.
mac	MAC authentication.
web-based	Web-based authentication.
convergence-endpoint	Convergence-endpoint authentication.

Default

300

Usage Guidelines

If no authentication type is specified, the idle timeout value is returned to 300 seconds for all authentication types.

Example

```
unconfigure netlogin idle-timeout mac
```

History

This command was first available in ExtremeXOS 16.1.

The **convergence-endpoint** option was added in ExtremeXOS 22.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure netlogin allowed-refresh-failures

```
unconfigure netlogin allowed-refresh-failures
```

Description

Restores the number refresh failures to the default value.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command allows you to restore the number of refresh failures allowed to the default value of 0.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure netlogin authentication database-order

```
unconfigure netlogin [mac | web-based] authentication database-order
```

Description

Restores the default order of database authentication protocols to use.

Syntax Description

mac	Specifies the MAC address to add.
mask	Specifies the number of bits to use for the mask.
default	Specifies the default entry.
encrypted	Used to display encrypted form of password in configuration files. Do not use.
password	Specifies the password to send for authentication.
ports	Specifies the port or port list to use for authentication.

Default

By default, the authentication order is *RADIUS*, local-user database.

Usage Guidelines

Use this command to restore the default configuration order for the database authentication protocols. For details, see the [Switch Engine 32.3 Feature License Requirements](#) document.

Example

The following command sets the database authentication order to RADIUS, local user database for MAC-based authentication:

```
unconfigure netlogin mac authentication database-order
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure netlogin authentication failure vlan

```
unconfigure netlogin authentication failure vlan vlan_name {ports
port_list}
```

<i>vlan_name</i>	Specifies the name of the authentication failure <u>VLAN</u> .
<i>port_list</i>	Specifies one or more ports or slots and ports. If the ports keyword is not used, the command applies to all ports.

Description

Disables authentication failure VLAN on network login enabled ports.

Syntax Description

Default

N/A.

Usage Guidelines

Use this command to disable authentication failure VLAN on network login enabled ports. When a supplicant fails authentication, it is moved to the authentication failure vlan and is given limited access until it passes the authentication.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure netlogin authentication service-unavailable vlan

```
unconfigure netlogin authentication service-unavailable vlan vlan_name
{ports port_list}
```

Description

Unconfigures authentication service unavailable VLAN on network login enabled ports.

Syntax Description

<i>vlan_name</i>	Specifies the name of the authentication service-unavailable VLAN.
<i>port_list</i>	Specifies one or more ports or slots and ports. If the ports keyword is not used, the command applies to all ports.

Default

Defaults to all network login enabled ports.

Usage Guidelines

This command unconfigures authentication service unavailable VLAN on the specified network login enabled ports. Authentication service unavailable VLAN is unconfigured on all the network login enabled ports, if no port is specifically mentioned.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure netlogin banner

```
unconfigure netlogin banner
```

Description

Unconfigures the network login page banner.

Syntax Description

This command has no arguments or variables.

Default

The default banner is the Extreme Networks logo.

Usage Guidelines

Use this command to unconfigure a netlogin banner.

After the command is issued, the configured banner specified is no longer displayed.

Example

The following command unconfigures the network login page banner:

```
unconfigure netlogin banner
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure netlogin dot1x guest-vlan

```
unconfigure netlogin dot1x guest-vlan {ports port_list | vlan_name}
```

Description

Unconfigures the guest *VLAN* feature for 802.1X authentication.

Syntax Description

<i>port_list</i>	Specifies one or more ports included in the guest VLAN.
<i>vlan_name</i>	Specifies all ports included in the guest VLAN.

Default

N/A.

Usage Guidelines

Use this command to unconfigure the guest VLAN for 802.1X authentication.

If you do not specify one or more ports or the VLAN name, this command unconfigures all of the 802.1X ports configured for the guest VLAN feature.

If you specify one or more ports, this command unconfigures the specified 802.1X ports for the guest VLAN feature.

If you specify the VLAN name, this command unconfigures all of the 802.1X ports configured for the specified guest VLAN.

Example

The following command unconfigures the guest VLAN feature for 802.1X authentication:

```
unconfigure netlogin dot1x guest-vlan
```

History

This command was first available in ExtremeXOS 11.2.

The ports option was added in ExtremeXOS 11.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure netlogin local-user security-profile

```
unconfigure netlogin local-user user-name security-profile
```

Description

Clears a previously associated security profile.

Syntax Description

<i>user-name</i>	Specifies the name of an existing local network login account.
------------------	--

Default

N/A.

Usage Guidelines

Use this command to clear any previously associated security profiles on the switch.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure netlogin ports

```
unconfigure netlogin ports [all | port_list] [allowed-users |  
authentication mode]
```

Description

Use this command to clear multiple authentication properties for one or more ports.

Syntax Description

all	Unconfigure all ports in the system.
<i>port_list</i>	List of ports to unconfigure.
allowed-users	Number of users allowed per port.
authentication	Unconfigure port authentication settings.
mode	Port authentication mode.

Default

N/A.

Usage Guidelines

None.

History

This command was first available in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure netlogin session-refresh

```
unconfigure netlogin session-refresh
```

Description

Restores the session refresh value to the default.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command allows you to restore the session refresh to the default value of 180 seconds.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure netlogin vlan

```
unconfigure netlogin vlan
```

Description

Unconfigures the VLAN for network login.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command unconfigures the VLAN used for unauthenticated clients. One VLAN needs to be configured per VR. To change the VLAN, network login needs to be disabled.

Example

The following command unconfigures the network login VLAN:

```
unconfigure netlogin vlan
```

History

This command was first available in ExtremeXOS 11.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure network-clock gtp ports

```
unconfigure network-clock gtp ports [port_list | all]
```

Description

Restores all configuration parameters on the specified ports to their default values. This command does not disable gPTP on the ports.

Syntax Description

<i>port_list</i>	Specifies one or more of the switch's physical ports.
all	Specifies all of the switch's physical ports.

Default

N/A.

Usage Guidelines

Use this command to restore all configuration parameters on the specified ports to their default values.

Example

```
# unconfigure network-clock gtp ports all
# unconfigure network-clock gtp ports 1,2
```

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the AVB feature pack license and have it installed. For complete information about software licensing, including how to obtain and upgrade your license, and which platforms support the AVB feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

unconfigure ospf

```
unconfigure ospf {vlan vlan-name | area area-identifier}
```

Description

Resets one or all *OSPF* interfaces to the default settings.

Syntax Description

<i>vlan-name</i>	Specifies a <i>VLAN</i> name.
<i>area-identifier</i>	Specifies an OSPF area.

Default

N/A.

Usage Guidelines

ExtremeXOS OSPF allows you to change certain configurable OSPF parameters on the fly. This command selectively resets the configurable parameters to their default values. Following is the list of parameters whose values will be reset to their defaults:

- Interface
 - Hello interval.
 - Dead interval.
 - Transmit delay.
 - Retransmit interval.
 - Priority.
 - Cost.
 - OSPF graceful restart helper mode.
- Area
 - All the parameters of interfaces associated with this area.
 - Inter-Area-Prefix_LSA Filter.
 - AS-External-LSA Filter.
- OSPF Global
 - All parameters of all areas in this OSPF domain.
 - SPF delay interval.
 - Interface cost metric table.
 - Route redistribution.
 - OSPF graceful restart.

Example

The following command resets the OSPF interface to the default settings on the VLAN accounting:

```
unconfigure ospf accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the **Switch Engine 32.3 Feature License Requirements** document.

unconfigure ospfv3

```
unconfigure ospfv3 {area area_identifier | vlan vlan_name | tunnel
                    tunnel_name}
```

Description

Resets one or all OSPFv3 interfaces to the default settings.

Syntax Description

<i>area_identifier</i>	Specifies an OSPFv3 area, a four-byte, dotted decimal number.
<i>vlan_name</i>	Specifies an IPv6 configured <u>VLAN</u> .
<i>tunnel_name</i>	Specifies an IPv6 tunnel.

Default

N/A.

Usage Guidelines

ExtremeXOS OSPFv3 allows you to change certain configurable OSPFv3 parameters on the fly. This command selectively resets the configurable parameters to their default values. The following is the list of parameters whose values will be reset to their defaults:

- Interface:
 - Hello Interval.
 - Dead Interval.
 - Transmit Delay.
 - Retransmit Interval.
 - Priority.
 - Cost.
- Area:
 - All the parameters of Interfaces associated with this area.
 - Inter-Area-Prefix-LSA Filter.
 - AS-External-LSA Filter.
- OSPFv3 Global:
 - SPF Delay interval.
 - Interface Cost metric Table.
 - Route Redistribution.

Example

The following command resets the OSPFv3 interface to the default settings on the VLAN accounting:

```
unconfigure ospfv3 accounting
```

The following command unconfigures the parameters of the area 0.0.0.1 (and all its associated interfaces):

```
unconfigure ospfv3 area 0.0.0.1
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with Base license, or higher, as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

unconfigure pim

```
unconfigure pim {ipv4 | ipv6} {vlan vlan_name | tunnel tunnel_name | ssm range | border}
```

Description

Resets all PIM settings on an IPv4 or IPv6 module, or on one or all [VLANs](#), to their default values.

Syntax Description

ipv4	Specifies the IPv4 module from which PIM is to be unconfigured.
ipv6	Specifies the IPv6 module from which PIM is to be unconfigured.
<i>vlan_name</i>	Specifies the VLAN from which PIM is to be unconfigured.
tunnel	Specifies the tunnel which PIM is to be unconfigured.
<i>tunnel_name</i>	Specifies the tunnel name.
border	Specifies the border.

Default

If no VLAN is specified, the configuration is reset for all PIM interfaces.

Usage Guidelines

If you unconfigure PIM globally using the `unconfigure pim {ipv4|ipv6}` command, you also unconfigure PIM-SSM, removing PIM-SSM range and SSM setting for PIM all interfaces. Static RP configuration is not removed in this case.

Example

The following command resets all PIM settings on the VLAN accounting:

```
unconfigure pim vlan accounting
```

History

This command was first available in ExtremeXOS 10.1.

The **ipv4** and **ipv6** keywords were added, giving an option to support this functionality in IPv6 as well, in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

unconfigure pim border

```
unconfigure pim {ipv4 | ipv6} {vlan} vlan_name border
```

Description

Unconfigures a PIM *VLAN* that has been configured as a border VLAN, which is used to demarcate a PIM domain.

Syntax Description

ipv4	Configures PIM functionality on IPv4 router interfaces.
ipv6	Configures PIM functionality on IPv6 router interfaces.
<i>vlan_name</i>	Specifies a VLAN name.

Default

By default, no PIM VLANs are configured as border VLANs.

Usage Guidelines

A PIM-SM domain is created by limiting the reach of PIM BSR advertisements. When a border VLAN is configured, PIM BSR advertisements are not forwarded out of the PIM VLAN. Use the `unconfigure pim border` command to remove the border functionality of the specified PIM VLAN.

Example

The following command unconfigures a PIM border on a VLAN called "vlan_border":

```
unconfigure pim vlan_border border
```


History

This command was first available in ExtremeXOS 12.0.

The IPv6 configuration option was added in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

unconfigure pim ssm range

```
unconfigure pim {ipv4 | ipv6} ssm range
```

Description

Unconfigures the range of multicast addresses for PIM SSM.

Syntax Description

ipv4	Configures PIM functionality on IPv4 router interfaces.
ipv6	Configures PIM functionality on IPv6 router interfaces.

Default

By default, no SSM range is configured.

Usage Guidelines

Initially, no range is configured for SSM. After a range is configured, you can remove the range with the `unconfigure pim ssm range` command.

When no range is configured for PIM SSM, the switch does not use PIM SSM for any multicast groups.

Example

The following command removes the PIM SSM range:

```
unconfigure pim ssm range
```

History

This command was first available in ExtremeXOS 11.4.

The **ipv4** and **ipv6** keywords were added, giving an option to support this functionality in IPv6 as well, in ExtremeXOS 15.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

NEW! `unconfigure pim tunnel`

```
unconfigure pim {ipv4}{tunnel tunnel_name}
```

Description

Resets all PIM tunnel settings on an IPv4 module.

Syntax Description

ipv4	Specifies the IPv4 module from which PIM is to be unconfigured.
tunnel	Specifies the tunnel which PIM is to be unconfigured.
<i>tunnel_name</i>	Specifies the tunnel name.

Default

If no VLAN is specified, the configuration is reset for all PIM interfaces.

Usage Guidelines

If you unconfigure PIM globally using the `unconfigure pim {ipv4|ipv6}` command, you also unconfigure PIM-SSM, removing PIM-SSM range and SSM setting for PIM all interfaces. Static RP configuration is not removed in this case.

Example

The following command resets all PIM settings on the tunnel accounting:

```
unconfigure pim tunnel accounting
```

History

This command was first available in ExtremeXOS 32.3.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the PIM feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

unconfigure policy all-rules

```
unconfigure policy all-rules
```

Description

Use this command to remove all admin and classification rules.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to remove all admin and classification rules.

History

This command was first available in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure policy app-signature group name

```
unconfigure policy app-signature group group [name name | custom]
```

Description

Deletes a user-defined policy application signature.

Syntax Description

app-signature	Configures application signature specific settings.
group	Configures application signature group-specific settings.
<i>group</i>	Specifies the group name.
name	Configures application signature display name-specific settings.
<i>name</i>	Specifies the display name assigned to the application signature.
custom	Removes all custom signature information from the specified group.

Default

N/A.

Usage Guidelines

The application signature groups are pre-defined (built-in), and additional ones cannot be created. There are pre-defined values for application signature names and patterns as well, and these cannot be modified or deleted.

Example

The following example removes all custom signature information from the group "E-commerce":

```
# unconfigure policy app-signature group "E-commerce" custom
```

History

This command was first available in ExtremeXOS 30.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure policy autoclear

```
unconfigure policy autoclear interval
```

Description

Disables the automatic clearing of rule usage statistics.

Syntax Description

autoclear	Designates setting parameters for auto-clearing the policy rule usage statistics.
interval	Designates disabling automatic clearing of rule usage statistics.

Default

N/A.

Usage Guidelines

If you have configured Syslog and/or trap actions to notify you when a policy rule is used by using the following command: `configure policy rule profile_index [{app-signature group group name name} | ether ether | icmp6type icmp6type | icmptype icmptype | ip6dest ip6dest | ipdestsocket ipdestsocket |`

```
ipfrag | ipproto ipproto | ipsourcesocket ipsourcesocket | iptos
iptos | ipttl ipttl | macdest macdest | macsource macsource | port
port | tcpdestportIP tcpdestportIP | tcpsourceportIP tcpsourceportIP |
udpdestportIP udpdestportIP | udpsourceportIP udpsourceportIP ] {mask
mask } {port-string [ port_string | all]} {storage-type [non-volatile
| volatile]} {drop | forward} {syslog syslog} {trap trap} {cos cos }
{mirror-destination control_index} {clear-mirror} , this command disables
automatic clearing of rule usage statistics.
```

To view the auto-clear interval, use the following command:

```
show policy autoclear interval
```

Example

The following example disables automatic clearing of rule usage statistics:

```
# unconfigure policy autoclear interval
# show policy autoclear interval
Stats autoclear interval DISABLED.
```

History

This command was first available in ExtremeXOS 30.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure policy captive-portal

```
unconfigure policy captive-portal web-redirect redirect_index server
server_id
```

Description

This command unconfigures a particular captive portal server. This disables the server from being used by an existing policy profile.

Syntax Description

web-redirect	Unconfigures web-redirect.
<i>redirect_index</i>	Specifies the web redirect index (range = 1-10).
server	Unconfigures a server for the web redirect index.
<i>server_id</i>	Specifies the server ID to disable.

Default

N/A

Example

The following example unconfigures a particular captive portal server (index 2) in web-redirect (index 1):

```
unconfigure policy captive-portal web-redirect 1 server 2
```

History

This command was first available in ExtremeXOS 22.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure policy captive-portal listening

```
unconfigure policy captive-portal listening [ socket_list | all ]
```

Description

This command removes previously defined L4 listening ports (sockets).

Syntax Description

listening	Configures captive portal HTTP listening ports (up to three L4 ports).
<i>socket_list</i>	List of L4 listening ports that you want removed, separated by commas (for example: 80,8080,2000).
all	Removes all listening L4 ports.

Default

N/A

Example

The following example removes captive portal L4 listening ports 80 and 8080:

```
# unconfigure policy captive-portal listening 80,8080
```

The following example removes all defined captive portal L4 listening ports:

```
# unconfigure policy captive-portal listening all
```

History

This command was first available in ExtremeXOS 22.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure policy convergence-endpoint all

```
unconfigure policy convergence-endpoint all
```

Description

This command unconfigures Convergence End Point (CEP) settings back to their factory default.

Syntax Description

This command has no arguments or variables.

Default

N/A

Example

The following example clears all CEP settings and restores the factory default:

```
# unconfigure policy convergence-endpoint all
```

History

This command was first available in ExtremeXOS 22.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure policy convergence-endpoint index

```
unconfigure policy convergence-endpoint index [cisco | lldp-med]
```

Description

This command unconfigures a global default policy index for a Convergence End Point (CEP) detection type.

Syntax Description

cisco	Unconfigures Cisco detection type only.
lldp-med	Unconfigures <u>LLDP</u> -MED detection type only.

Default

N/A

Example

The following example removes the default policy for Cisco type CEPs:

```
# unconfigure policy convergence-endpoint index cisco
```

History

This command was first available in ExtremeXOS 22.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure policy invalid action

```
configure policy invalid action
```

Description

This command removes the actions configured to be taken for an invalid policy.

Default

N/A

Example

The following command removes the actions configured to be taken for an invalid policy:

```
configure policy invalid action
```

History

This command was first available in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure policy mactable

```
unconfigure policy mactable [ response | vlan_list]
```

Description

Use this command to clear a VLAN to policy mapping table entry or to reset the mactable response to the default value of policy mode.

Syntax Description

<code>mactable</code>	Clear VLAN ID - Policy Profile mappings table.
<code>response</code>	Specifies that the response should be reset to the default value of policy.
<code>vlan_list</code>	Specifies the VLAN ID or range of IDs (1 to 4094) to clear.

Default

N/A.

Example

This example resets the policy mappings table response to the default value of policy.

```
unconfigure policy mactable response
```

History

This command was first available in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure policy profile

```
unconfigure policy profile [all | profile_index]
```

Description

Use this command to delete a policy profile entry.

Syntax Description

<i>profile</i>	Clear policy profile entries.
all	Clear all policy profiles.
<i>profile_index</i>	Specifies the index number of the policy profile entry to be deleted. Valid values are 1 to 1023.

Default

N/A.

Example

This example shows how to delete policy profile 8:

```
unconfigure policy profile 8
```

History

This command was first available in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure policy rule

```
unconfigure policy rule [ profile_index ] [ all-pid-entries ] | [[ ether
| icmp6type | icmptype | ip6dest | ipdestsocket | ipfrag | ipproto
| ipsourcesocket | iptos | ipttl | macdest | macsource | port |
tcpsourceportIP | udpsourceportIP | tcpdestportIP | udpdestportIP ]
{ app-signature } [ all-traffic-entries | data ] { mask mask } { port-string
port_string | all } }
```

Description

Use this command to delete one or all policy classification rule entries.

Syntax Description

See [configure policy rule](#) on page 1127 for a full list of syntax descriptions.

all-pid-entries	Clear all entries associated with this Policy ID.
all-traffic-entries	Clear all entries associated with this traffic rule.
<i>data</i>	Data corresponding to rule-type option.
app-signature	Clears configured application signatures.

Default

When applicable, data, mask, and port-string must be specified for individual rules to be cleared.

Example

This example shows how to delete all classification rule entries associated with policy profile 1 from all ports:

```
unconfigure policy rule 1 all-pid-entries
```

History

This command was first available in ExtremeXOS 16.1.

ICMP and ICMPv6 rule types added in ExtremeXOS 22.5.

The **app-signature** option was added in ExtremeXOS 30.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure policy syslog

```
unconfigure policy syslog {machine-readable} {extended-format} {every-time}
```

Description

Resets Syslog parameters for policy rules to their default settings.

Syntax Description

syslog	Designates unconfiguring selected Syslog parameters for policy rules.
machine-readable	Clears the control for device formatting of rule usage messages to the default setting of disabled (decimal format).
extended-format	Clears the control for extended rule usage message formatting to the default setting of disabled (not extended).
every-time	Sets the Syslog message sent on rule usage to the default of on first usage only (disabled).

Default

By default, Syslog messages are only sent on first use of a rule.

By default, **extended-format** and **machine-readable** are disabled (not extended and in decimal format).

Usage Guidelines

This command allows you to reset parameters to their defaults for Syslog messages that are sent when a policy rule is used as set up in the command `configure policy rule profile_index [{app-signature group group name name} | ether ether | icmp6type icmp6type | icmp type icmp type | ip6dest ip6dest | ipdestsocket ipdestsocket | ipfrag | ipproto ipproto | ipsourcesocket ipsourcesocket | iptos iptos | ipttl ipttl | macdest macdest | macsource macsource | port port | tcpdestportIP tcpdestportIP | tcpsourceportIP tcpsourceportIP | udpdestportIP udpdestportIP | udpsourceportIP udpsourceportIP] {mask mask } {port-string [port_string | all]} {storage-type [non-volatile | volatile]} {drop | forward} {syslog syslog} {trap trap} {cos cos } {mirror-destination control_index} {clear-mirror} .`

Example

The following clears the setting for when Syslog messages are sent when a rule is used to the default of only on first rule usage:

```
#unconfigure policy syslog every-time
```

History

This command was first available in ExtremeXOS 30.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure policy vlanauthorization

```
unconfigure policy vlanauthorization [{port port_list } | all ]
```

Description

This command clears VLAN Authorization to the default state for a specific port, list of ports, or all ports.

Syntax Description

port	Specifies ports to clear VLAN Authorization to default state.
<i>port_list</i>	List of ports to clear to default state.
all	Clear all VLAN Authorization state to default.

Default

N/A.

Usage Guidelines

None.

Example

This example shows how to clear VLAN Authorization for all ports:

```
x450G2-48t-10G4.3 # unconfigure policy vlanauthorization all
```

This example shows how to clear VLAN Authorization for ports 1:1-4:

```
x450G2-48t-10G4.3 # unconfigure policy vlanauthorization port 1:1-4
```

History

This command was first available in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure port description-string

```
unconfigure ports port_list description string
```

Description

Unconfigures a description string setting.

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
------------------	---

Default

N/A.

Usage Guidelines

Use this command to unconfigure a port description.

Example

The following command unconfigures the port description string:

```
unconfigure ports 1:3
```

History

This command was available in ExtremeXOS 15.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure ports display string

```
unconfigure ports port_list display-string
```

Description

Clears the user-defined display string from one or more ports.

Syntax Description

<i>port_list</i>	Specifies one or more ports or slots and ports.
------------------	---

Default

N/A.

Usage Guidelines

This command removes the display string that you configured using the [configure ports display-string](#) command.

Example

The following command clears the user-defined display string from port 4 on a switch:

```
unconfigure ports 4 display-string
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure ports link-flap-detection

```
unconfigure ports [port_list | all] link-flap-detection {interval}  
                {threshold} {disable-time}
```

Description

Sets interval, threshold (maximum number of link down events), and disable time to defaults for link-flap detection.

Syntax Description

ports	Physical ports.
<i>port_list</i>	List of ports to set link-flap detection characteristics back to defaults upon.
all	Sets link-flap detection characteristics to defaults on all ports in the system.
link-flap-detection	Specifies link-flap detection.
interval	Resets the time interval for collecting link-flap events to the default of 5 seconds.
threshold	Resets the number of link-flap events tolerated before action is taken to the default of 10.
disable-time	Resets the time period a port remains disabled after excessive link flapping to the default of 300 seconds.

Default

N/A

Example

The following example resets the link-flap detection interval and threshold to their defaults on ports 1-4.

```
unconfigure ports 1-4 link-flap-detection interval threshold
```

History

This command was first available in ExtremeXOS 22.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure ports monitor vlan

```
unconfigure ports [port_list | all] monitor vlan [ vlan_name |  
    vlan_list ]
```

Description

Stops counting VLAN statistics on a port or group of ports.

Syntax Description

<i>port_list</i>	Specifies one or more ports. May be in the form: 1, 2, 3-5, 2:5, 2:6-2:8.
<i>vlan_name</i>	Specifies a VLAN name.
<i>vlan_list</i>	Specifies a list of VLANs.

Default

N/A.

Usage Guidelines

None.

Example

The following command removes monitoring for ports on a supported switch on the VLAN named accounting:

```
# unconfigure ports 1-6 monitor vlan accounting
```

History

This command was first available in ExtremeXOS 12.0.

Support for SummitStack and ExtremeSwitching series switches was added in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure ports redundant

```
unconfigure ports port_list redundant
```

Description

Clears a previously configured software-controlled redundant port.

Syntax Description

<i>port_list</i>	This refers to the primary port of the redundant pair and specifies one or more ports or slots and ports.
------------------	---

Default

N/A.

Usage Guidelines

The list of port numbers or the port display string specifies the primary port(s).

Example

The following command unconfigures a software-controlled redundant port on a switch:

```
unconfigure ports 3 redundant
```

History

This command was available in ExtremeXOS 11.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure process group

```
unconfigure process group
```

Description

This command restores the default settings for memory and CPU limits for both the "Vital" and "Other" process groups.

Syntax Description

N/A.

Default

None.

Usage Guidelines

If one of the process groups is already consuming more memory than the default limit, an error may appear and the command may not execute successfully:

```
Warning: Unable to restore memory limits to default values. One or both of the groups  
are already consuming more memory than their default limits (95% for "EXOS" and 5%  
for "Other"). CPU limits have been restored to their default values.
```

Example

The following example restores the default memory and CPU limits of the process control groups:

```
unconfigure process group
```

History

This command was first available in ExtremeXOS 22.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure qosprofile

unconfigure qosprofile {**ports** [*port_list*|**all**]}

Description

Returns the rate-shaping parameters for all QoS profiles on the specified ports to the default values.

Syntax Description

<i>port_list</i>	Specifies the ports on which to unconfigure QoS profiles.
all	Specifies that this command applies to all ports on the device.

Default

The default values for egress bandwidth on all supported platforms are:

- Minimum bandwidth—0%
- Maximum bandwidth—100%

The default values for egress priority and ingress QoS profiles differ by platform as described in the following sections.

The default values for the two default egress QoS profiles (QP1 and QP8) are:

- Maximum buffer—100% (as set by the [configure qosprofile](#) command)
- Maximum buffer override—100% (as set by the [configure qosprofile](#) command)
- Weight—1
- WRED—See the [configure qosprofile wred](#) command description.

Usage Guidelines

None.

Example

The following command resets the QoS profiles for all ports to default settings:

```
unconfigure qosprofile
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

The egress and ports parameters are available only on all platforms.

unconfigure qosprofile wred

```
unconfigure qosprofile wred {ports [port_list | all]}
```

Description

Removes the WRED configuration for all QoS profiles on the specified port or all ports.

Syntax Description

<i>port_list</i>	Specifies a list of slots and ports from which the WRED configuration is removed. Specify ports in the following formats: 3-5, 2:5, 2:6-2:8.
all	Specifies that this command applies to all ports on the device.

Default

N/A.

Usage Guidelines

None.

Example

The following example removes the WRED configuration for port 3:

```
# unconfigure qosprofile wred port 3
```

History

This command was first available in ExtremeXOS 12.7.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure qosscheduler ports

```
unconfigure qosscheduler ports [port_list | port_group | all]
```

Description

When a port or PortGroup has not been configured for per-port scheduling, it uses the global scheduling algorithm. If a port or PortGroup has been configured with per-port scheduling, this command removes the per-port scheduling and the port/PortGroup will use the global scheduling algorithm.

Syntax Description

<i>port_list</i>	Port list.
<i>port_group</i>	Port group name.
all	All ports.

Default

N/A.

Usage Guidelines

None.

History

This command was first available in ExtremeXOS 16.1.

Platform Availability

All platforms.

unconfigure radius

```
unconfigure radius {mgmt-access | netlogin} {server [primary |  
                  secondary] }
```

Description

Unconfigures the *RADIUS* client configuration.

Syntax Description

mgmt-access	Specifies the switch management RADIUS authentication server.
netlogin	Specifies the network login RADIUS authentication server.

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

Default

Unconfigures both primary and secondary servers for management and network login.

Usage Guidelines

If you do not specify any keywords, this command unconfigures both the primary and secondary servers for management and network login.

The following list describes the available keywords:

- **mgmt-access**—Use this keyword to unconfigure only the server(s) for management functions.
- **netlogin**—Use this keyword to unconfigure only the server(s) for network login.
- **primary**—Use this keyword to specify only the primary RADIUS sever.
- **secondary**—Use this keyword to specify only the secondary RADIUS server.

Example

The following command unconfigures the secondary RADIUS server settings for both management and network login:

```
unconfigure radius server secondary
```

The following command unconfigures the secondary RADIUS server settings for only network login:

```
unconfigure radius netlogin server secondary
```

The following command unconfigures all RADIUS server settings for only management functions:

```
unconfigure radius mgmt-access
```

History

This command was first available in ExtremeXOS 10.1.

The **mgmt-access** and **netlogin** keywords were added in ExtremeXOS 11.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure radius-accounting

```
unconfigure radius-accounting {mgmt-access | netlogin} {server [primary  
| secondary] }
```

Description

Unconfigures the *RADIUS* accounting server configuration.

Syntax Description

mgmt-access	Specifies the switch management RADIUS accounting server.
netlogin	Specifies the network login RADIUS accounting server.
primary	Unconfigures the primary RADIUS accounting server.
secondary	Unconfigures the secondary RADIUS accounting server.

Default

Unconfigures both the primary and secondary accounting servers for management and network login.

Usage Guidelines

If you do not specify any keywords, this command unconfigures both the primary and secondary accounting servers for management and network login.

The following list describes the available keywords:

- **mgmt-access**—Use this keyword to unconfigure only the accounting server(s) for management functions.
- **netlogin**—Use this keyword to unconfigure only the accounting server(s) for network login.
- **primary**—Use this keyword to specify only the primary RADIUS accounting sever.
- **secondary**—Use this keyword to specify only the secondary RADIUS accounting server.

Example

The following command unconfigures the secondary RADIUS accounting server settings for both management and network login:

```
unconfigure radius-accounting server secondary
```

The following command unconfigures the secondary RADIUS accounting server settings for only network login:

```
unconfigure radius-accounting netlogin server secondary
```

The following command unconfigures all RADIUS accounting server settings for only management functions:

```
unconfigure radius-accounting mgmt-access
```

History

This command was first available in ExtremeXOS 10.1.

The mgmt-access and netlogin keywords were added in ExtremeXOS 11.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure radius-accounting server

```
unconfigure radius-accounting [server index]
```

Description

This command removes the *RADIUS* server configuration for a given server. Having *index* be in its own setting is done to remove the possibility to set the mgmt-access or netlogin setting prior to index.

Syntax Description

server	RADIUS accounting server.
<i>index</i>	RADIUS accounting server index.

Default

N/A.

Usage Guidelines

None.

History

This command was first available in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure radius server

```
unconfigure radius {dynamic-authorization [server index]
```

Description

This command removes the *RADIUS* server configuration for a given server. Having *index* be in its own setting is done to remove the possibility to set the mgmt-access or netlogin setting prior to index.

Syntax Description

server	RADIUS server.
<i>index</i>	RADIUS server index.
dynamic-authorization	Specifies dynamic-authorization.

Default

N/A.

Usage Guidelines

None.

History

This command was first available in ExtremeXOS 16.1.

The **dynamic-authorization** keyword was added in ExtremeXOS 22.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure rip

```
unconfigure rip {vlan vlan-name | all}
```

Description

Resets all *RIP* parameters to the default for all *VLANs* or for the specified VLAN.

Syntax Description

<i>vlan-name</i>	Specifies a VLAN name.
------------------	------------------------

Default

All.

Usage Guidelines

Does not change the enable/disable state of the RIP settings.

Example

The following command resets the RIP configuration to the default for the VLAN finance:

```
# unconfigure rip finance
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on all platforms with a Base license or higher as described in the [*Switch Engine 32.3 Feature License Requirements*](#) document.

unconfigure ripng

```
unconfigure ripng {vlan vlan-name | tunnel tunnel-name | vlan all |  
                  tunnel all}
```

Description

Resets *RIPng* parameters to the default value.

Syntax Description

<i>vlan-name</i>	Specifies an IPv6 configured <i>VLAN</i> .
<i>tunnel-name</i>	Specifies an IPv6 tunnel.
all	Specifies either all IPv6 configured VLANs or all IPv6 tunnels.

Default

N/A.

Usage Guidelines

Issuing the command unconfigure ripng resets all the interfaces and the global configuration to the defaults, and disables RIPng, as that is the default.

Example

The following command resets the RIPng configuration to the default for the VLAN finance:

```
unconfigure rip finance
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on platforms with an Edge, Advanced Edge, or Core license. For licensing information, see the [Switch Engine 32.3 Feature License Requirements](#) document.

unconfigure sflow

```
unconfigure sflow
```

Description

Resets all the sFlow values to the default values.

Syntax Description

This command has no arguments or variables.

Default

The default values for sFlow are as follows:

- sFlow agent IP address—0.0.0.0.
- sampling frequency—sample one every 8192 packets.
- polling interval—20 seconds.
- maximum CPU sample limit—2000 samples per second.

sFlow is unconfigured and disabled on all ports.

Usage Guidelines

This command resets sFlow values to the default values, and removes any port configurations, and any sFlow collectors configured on the switch.

Example

The following command unconfigures sFlow:

```
unconfigure sflow
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure sflow agent

```
unconfigure sflow agent
```

Description

Resets the sFlow agent's IP address to the default value.

Syntax Description

This command has no arguments or variables.

Default

The default IP address is 0.0.0.0.

Usage Guidelines

This command resets the sFlow agent IP address to its default value.

Example

The following command resets the agent IP back to the management IP address:

```
unconfigure sflow agent
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure sflow collector

```
unconfigure sflow collector {ipaddress} ipaddress {port udp-port-number}  
{vr vr_name}
```

Description

Unconfigures the sFlow collector.

Syntax Description

<i>ipaddress</i>	Specifies the IP address of the collector to reset.
<i>udp-port-number</i>	Specifies the UDP port.
<i>vr_name</i>	Specifies which virtual router. Note: User-created VRs are supported only on the platforms listed for this feature in the Switch Engine 32.3 Feature License Requirements document.

Default

The following values are the defaults for this command:

- UDP port number—6343
- Virtual router—*VR-Mgmt* (previously called VR-0)

Usage Guidelines

This command allows you to reset the specified sFlow collector parameters to the default values.

Both the commands `unconfigure ports monitor vlan` and `unconfigure sflow collector` will reset the collector parameters to the default.

Example

The following command removes the collector at IP address 192.168.57.1:

```
unconfigure sflow collector ipaddress 192.168.57.1
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure sflow ports

```
unconfigure sflow ports port_list
```

Description

Removes the specified ports from the sFlow configuration, and stops sampling them.

Syntax Description

<code>port_list</code>	Specifies one or more ports or slots and ports.
------------------------	---

Default

N/A.

Usage Guidelines

This command removes the specified ports from the sFlow configuration, and stops sampling them.

Example

The following command unconfigures sFlow on the ports 2:5-2:7:

```
unconfigure sflow ports 2:5-2:7
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure slot

```
unconfigure slot slot
```

Description

Clears a slot of a previously assigned switch type.

Syntax Description

<code>slot</code>	Specifies a slot (node) on a SummitStack.
-------------------	---

Default

N/A.

Usage Guidelines

If you issue the `unconfigure ports wan-phy` command on a slot containing a switch with any ports configured for software-controlled redundancy, this command wipes away all software-controlled

redundancy on both ports; both ports return to normal. Refer to the [Switch Engine 32.3 User Guide](#) for more information on software-controlled redundant ports.

Example

The following command clears node (slot) 4 of a previously assigned switch type:

```
unconfigure slot 4
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available only on SummitStack.

unconfigure ssl certificate

```
unconfigure ssl certificate [trusted-ca | ocsp-signature-ca] [file_name  
| all ]
```

Description

Removes a trusted CA certificate or OCSP response signature CA certificate from the switch.

Syntax Description

trusted-ca	Specifies trusted CA certificates. Removes the CA certificate file(s) from folder <code>/config/ssl/trusted_ca_certs/</code> .
ocsp-signature-ca	Specifies OCSP signature CA certificates. Removes the signature CA file(s) from folder <code>/config/ssl/signature_ca_certs/</code> .
<i>file_name</i>	Specifies a single CA certificate file for removal.
all	Specifies all CA certificate files for removal.

Default

None.

Example

The following example removes all trusted CA certificates from the switch:

```
unconfigure ssl certificate trusted-ca all
```

History

This command was first available in ExtremeXOS 22.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure stacking

```
unconfigure stacking {node-address node_address | slot slot_number}
```

Description

This command resets most stacking parameters to the default or unconfigured values.

Syntax Description

<i>node_address</i>	Specifies the MAC address of a node in the stack. To view the MAC addresses for all nodes in a stack, enter the show stacking command.
<i>slot_number</i>	Specifies the slot number of the target node. To view the slot numbers, enter the show stacking command.

Default

N/A.

Usage Guidelines

Run this command from any node. If you do not specify a target node, the stacking parameters on all nodes are reset.

This command resets the stacking parameters shown in the following table.

Table 48: Stacking Configuration Items, Time of Effect and Default Value

Configuration Item	Takes Effect	Default Value
Stacking Mode	At boot time	Disabled
Slot Number	At boot time	1
Master-Capable	At boot time	Yes
License Restriction	At boot time	Not Configured
Priority	At next master election	Automatic
Alternate IP Address	Immediately	Not Configured
Stack MAC	At boot time	Not Configured

This command does not reset the stacking parameters configured with the following commands that use the stacking-support keyword:

- `configure stacking-support`
- `disable stacking-support`
- `enable stacking-support`

Example

To unconfigure the stacking parameters of all nodes in the stack topology:

```
unconfigure stacking
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available with all licenses and platforms that support the SummitStack feature. For information about which licenses and platforms support the SummitStack feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

unconfigure stacking alternate-ip-address

```
unconfigure stacking {node-address node_address | slot slot_number}  
alternate-ip-address
```

Description

Removes the configured alternate management IP address from the specified node.

If no node is specified, the alternate management IP address is removed from every node. The change takes effect immediately for all nodes operating in stacking mode.

Syntax Description

<i>node_address</i>	Specifies the MAC address of a node in the stack. To view the MAC addresses for all nodes in a stack, enter the <code>show stacking</code> command.
<i>slot_number</i>	Specifies the slot number of the target active node. To view the slot numbers, enter the <code>show stacking</code> command.

Default

N/A.

Usage Guidelines

Run this command from any node.

Example

To unconfigure stacking alternate-ip-address on a node:

```
unconfigure stacking node-address 00:04:96:26:6b:ed alternate-ip-address
```

To unconfigure the stacking alternate IP address configured on the active node in slot 4:

```
unconfigure stacking slot 4 alternate-ip-address
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available with all licenses and platforms that support the SummitStack feature. For information about which licenses and platforms support the SummitStack feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

unconfigure stacking license-level

```
unconfigure stacking {node-address node_address | slot slot_number}  
    license-level
```

Description

This command removes a previously configured license level restriction.

Syntax Description

<i>node_address</i>	Specifies the MAC address of a node in the stack. To view the MAC addresses for all nodes in a stack, enter the show stacking command.
<i>slot_number</i>	Specifies the slot number of the target node. To view the slot numbers, enter the show stacking command.

Default

N/A.

Usage Guidelines

If no node is specified, the licensing restriction is removed from all nodes in the stack topology.

After the command is executed, the following message appears:

```
This command will take effect at the next reboot of the specified node(s).
```

Example

To unconfigure the stacking license level on a node:

```
unconfigure stacking node-address 00:04:96:26:6b:ed license-level
```

To unconfigure the stacking license level configured on slot 4:

```
unconfigure stacking slot 4 license-level
```

History

This command was first available in ExtremeXOS 12.0.

Platform Availability

This command is available with all licenses and platforms that support the SummitStack feature. For information about which licenses and platforms support the SummitStack feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

unconfigure stacking-support

```
unconfigure stacking-support
```

Description

This command resets the stacking parameters configured with commands that use the stacking-support keyword.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Run this command from the local node on which you want to reset stacking-support parameters.

This command resets the stacking parameters configured with the following commands that use the stacking-support keyword:

- configure stacking-support

- `disable stacking-support`
- `enable stacking-support`

Example

To unconfigure the stacking-support parameters on the local node, use the following command:

```
unconfigure stacking-support
The stacking-support configuration has been reset.
The defaults will take effect at the next reboot of this switch.
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on all platforms that support alternate stack port selection or permit disabling of the stacking-support option.

unconfigure stpd ports link-type

```
unconfigure stpd stpd_name ports link-type port_list
```

Description

Returns the specified port to the factory default setting of broadcast link.

Syntax Description

<i>stpd_name</i>	Specifies an <i>STPD</i> name on the switch.
<i>port_list</i>	Specifies one or more ports or slots and ports.

Default

All ports are broadcast link types.

Usage Guidelines

If your STPD has the same name as another component, for example a *VLAN*, you must enter the stpd keyword to specify the STPD. If your STPD has a name unique only to that STPD, the keyword stpd is optional.

If the switch operates in 802.1D mode, any configured port link type will behave the same as the broadcast link type.

In an *MSTP* environment, configure the same link types for the CIST and all MSTIs.

Example

The following command configures slot 2, ports 1 through 4 to return to the factory default of broadcast links in STPD s1:

```
unconfigure stpd s1 ports link-type 2:1-2:4
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure stpd

```
unconfigure stpd {stpd_name}
```

Description

Restores default *STP* values to a particular *STPD* or all STPDs.

Syntax Description

<i>stpd_name</i>	Specifies an STPD name on the switch.
------------------	---------------------------------------

Default

N/A.

Usage Guidelines

If you create an STPD with a unique name, the keyword stpd is optional.

Use this command to restore default STP values to a particular STPD. If you want to restore default STP values on all STPDs, do not specify a spanning tree name.

Example

The following command restores default values to an STPD named Backbone_st:

```
unconfigure stpd backbone_st
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure switch

```
unconfigure switch {all | erase [all | nvram]}
```

Description

Returns the switch configuration to its factory default settings and reboots the switch.

Syntax Description

all	Specifies that the entire configuration should be changed to the default values, including the management IP address, failsafe account, and SummitStack-specific parameters, and the switch rebooted.
erase all	All data such as loaded ExtremeXOS images (both partitions), configuration files, policy files, non-volatile memory content, and switch settings are overwritten. This renders the switch inoperable until you perform a bootrom rescue. The operation reboots the system automatically to complete the erasure. The entire operation may take up to 20 minutes (it could take longer if an SSD device is attached). Do not disconnect the power during this time.
erase nvr am	Data in non-volatile memory such as selected configuration, selection image partition, log messages are overwritten. Switch boots up with primary image. Any unsaved configuration changes are lost and the switch reboots.

Default

N/A.

Usage Guidelines

Use `unconfigure switch` to reset the configuration to factory defaults, but without erasing the configuration. This preserves users account information, date and time settings, SummitStack configuration, and so on.

Include the parameter **all** to clear the entire current configuration, including all switch and SummitStack parameters, and reboot using the last used image and factory default configuration.

The command `unconfigure switch all` does not clear licensing information. The license cannot be disabled after it is enabled on the switch.

To remove all virtual machines (VMs) and their installation files, use the command with the **all** option:



Note

If you do not use the **all** option, VMs/install files are not removed.

For SummitStack only.

The **all** option also resets all stacking-specific parameters to defaults. To reset only the stacking-specific parameters to defaults, enter the `unconfigure stacking` command.

Beginning with ExtremeXOS 12.5, stacking support and stacking port selection are reset only on the local node. When stacking support of any kind is supported on the platform, the following message is added to the output that is shown on the console after this command has been confirmed:

```
Stacking-support will be unconfigured on this node only.
```

Example

The following command preserves the entire current configuration (but does not reload the current configuration after the switch reboots) and reboots the switch or SummitStack using the last specified saved image and factory default configuration:

```
unconfigure switch all
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure tacacs

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

Description

Unconfigures the TACACS+ server 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 settings:

```
unconfigure tacacs
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure tacacs-accounting

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

Description

Unconfigures the TACACS+ accounting server 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 settings:

```
unconfigure tacacs-accounting
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure timezone

```
unconfigure timezone
```

Description

Reverts the timezone back to its default value.

Syntax Description

timezone	Sets the timezone to the default value (UTC 0 noautodst).
-----------------	---

Default

The default value of UTC 0 noautodst.

Usage Guidelines

None.

Example

The following command reverts the timezone back to its default value:

```
unconfigure timezone
```

History

This command was first available in ExtremeXOS 31.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure trusted-ports trust-for dhcp-server

```
unconfigure trusted-ports [ports | all] trust-for dhcp-server
```

Description

Unconfigures, disables one or more *DHCP* trusted ports.

Syntax Description

<i>ports</i>	Specifies one or more trusted ports.
all	Specifies all trusted ports.

Default

N/A.

Usage Guidelines

Use this command to disable one or more DHCP trusted ports.

Displaying DHCP Trusted Server Information

To display the DHCP snooping configuration settings, including DHCP trusted ports if configured, use the following command:

```
show ip-security dhcp-snooping {vlan} vlan_name
```

To display any violations that occur, including those on DHCP trusted ports if configured, use the following command:

```
show ip-security dhcp-snooping violations {vlan} vlan_name
```

Example

The following command unconfigures ports 2:2 and 2:3 as trusted ports:

```
unconfigure trusted-ports 2:2-2:3 trust-for dhcp-server
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure tunnel

```
unconfigure tunnel tunnel_name ipaddress ipv6_address_mask
```

Description

Unconfigures an IPv6 address/prefix route from a tunnel.

Syntax Description

<i>tunnel_name</i>	Specifies an IPv6 tunnel.
<i>ipv6_address_mask</i>	Specifies an IPv6 address / IPv6 prefix length.

Default

N/A.

Usage Guidelines

Use this command to unconfigure an IPv6 address/prefix route from the specified tunnel.

Example

The following example unconfigures the 6in4 tunnel "link39" with the address 2001:db8::1111/64

```
unconfigure tunnel link39 2001:db8::1111/64
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 interworking feature in the [Switch Engine 32.3 Feature License Requirements](#) document.ss="- topic/ph "/>

unconfigure upm event

```
unconfigure upm event upm-event profile profile-name ports port_list
```

Description

Unconfigures the event from the specified profile and port list.

Syntax Description

<i>upm-event</i>	Specifies the type of event to be unconfigured.
<i>profile-name</i>	Specifies the profile from which the event is unconfigured.
<i>port-list</i>	Unconfigures the UPM profile from the specified port list.

Default

N/A.

Usage Guidelines

This command removes an event from the specified profile and port list.

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the Universal Port feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

unconfigure upm timer

```
unconfigure upm timer timer-name profile profile-name
```

Description

Removes a UPM profile from a UPM timer.

Syntax Description

<i>timer-name</i>	Unconfigures the specified UPM timer and deactivates any running timer.
<i>profile-name</i>	Removes the specified profile from the UPM timer.

Default

N/A.

Usage Guidelines

Use this command to unconfigure a timer setting. This command does not delete the timer.



Note

The specified timer is stopped by this command, even if it has been activated.

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on platforms that support the appropriate license. For complete information about software licensing, including how to obtain and upgrade your license and which licenses support the Universal Port feature, see the [Switch Engine 32.3 Feature License Requirements](#) document.

unconfigure vlan description

```
unconfigure {vlan} vlan_name description
```

Description

Removes the description for the specified VLAN.

Syntax Description

<i>vlan_name</i>	Specifies the VLAN name.
------------------	--------------------------

Default

N/A.

Usage Guidelines

None.

Example

The following example removes the description from VLAN vlan1:

```
unconfigure vlan vlan1 description
```

History

This command was first available in ExtremeXOS 12.4.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure vlan dhcp

```
unconfigure vlan vlan_name dhcp
```

Description

Unconfigure all the DHCP configuration information for the specified VLAN.

Syntax Description

<i>vlan_name</i>	Specifies the VLAN on which to unconfigure DHCP.
------------------	--

Default

N/A.

Usage Guidelines

None.

Example

The following command unconfigures the DHCP server for the VLAN temporary:

```
unconfigure temporary dhcp
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure vlan dhcp-address-range

```
unconfigure vlan vlan_name dhcp-address-range
```

Description

Unconfigure the *DHCP* address range information for the specified *VLAN*.

Syntax Description

<i>vlan_name</i>	Specifies the VLAN on which to unconfigure DHCP.
------------------	--

Default

N/A.

Usage Guidelines

None.

Example

The following command unconfigures the DHCP address range for the VLAN temporary:

```
unconfigure temporary dhcp-address-range
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure vlan dhcp-options

```
unconfigure {vlan} vlan_name dhcp-options {[ default-gateway | dns-  
server {primary | secondary} | wins-server]}
```

Description

Unconfigure the *DHCP* option information for the specified *VLAN*.

Syntax Description

<i>vlan_name</i>	Specifies the VLAN on which to unconfigure DHCP.
default-gateway	Specifies the router option.
dns-server	Specifies the Domain Name Server (DNS) option.
primary	Specifies the primary DNS option.
secondary	Specifies the secondary DNS option.
wins-server	Specifies the NetBIOS name server (NBNS) option.

Default

N/A.

Usage Guidelines

None.

Example

The following command unconfigures the DHCP options for the VLAN temporary:

```
unconfigure temporary dhcp-options
```

History

This command was first available in ExtremeXOS 11.0.

The primary and secondary DNS options were added in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure vlan ipaddress

```
unconfigure [ {vlan} vlan_name | vlan vlan_list] ipaddress
```

Description

Removes the IP address of the VLAN or a VMAN. With no parameters, the command removes the primary IPv4 address on the specified VLAN. Using the IPv6 parameters, you can remove specified IPv6 addresses from the specified VLAN.

Syntax Description

<i>vlan_name</i>	Specifies a VLAN or VMAN name.
<i>vlan_list</i>	Specifies a VLAN list of IDs.
<i>ipv6_address_mask</i>	Specifies an IPv6 address using the format of IPv6-address/prefix-length, where IPv6 is the 128-bit address and the prefix length specifies the number of leftmost bits that comprise the prefix.

Default

Removes the primary IPv4 address from the specified VLAN or VMAN.

Usage Guidelines



Note

With IPv6, you cannot remove the last link local IPv6 address until all global IPv6 addresses are removed. For configurations, you cannot remove an IP address from a VLAN until after you delete the MLAG peer.

Example

The following command removes the primary IPv4 address from the VLAN "accounting":

```
unconfigure vlan accounting ipaddress
```

The following command removes an IPv6 addresses from the VLAN "finance":

```
unconfigure vlan finance ipaddress 3ffe::1
```

History

This command was first available in ExtremeXOS 10.1.

The IPv6 parameters were added in ExtremeXOS 11.2.

The `vlan_list` option was added in ExtremeXOS 16.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure vlan router-discovery

```
unconfigure vlan vlan_name router-discovery {ipv6}
```

Description

Unconfigures all the router-discovery parameters and resets them to their respective default values.

Syntax Description

<i>vlan_name</i>	Specifies an IPv6 configured <u>VLAN</u> .
------------------	--

Default

N/A.

Usage Guidelines

Each of the router-discovery parameters is set to the default value. For example, the default-lifetime parameter is set to 1800 seconds. The default value for each of the router-discovery parameters is listed in the corresponding `configure vlan router-discovery` command description.

Example

The following example unconfigures all the router-discovery parameters for the VLAN `top_floor`:

```
unconfigure vlan top_floor router-discovery
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

unconfigure vlan router-discovery default-lifetime

```
unconfigure vlan vlan_name router-discovery {ipv6} default-lifetime
```

Description

Unconfigures the router lifetime value sent in router discovery advertisements on the VLAN.

Syntax Description

<i>vlan_name</i>	Specifies an IPv6 configured VLAN.
------------------	------------------------------------

Default

N/A.

Usage Guidelines

This command sets the default-lifetime parameter to the default value of 1800 seconds.

Example

The following command unconfigures the default-lifetime for the VLAN "top_floor":

```
unconfigure vlan top_floor router-discovery default-lifetime
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

unconfigure vlan router-discovery hop-limit

```
unconfigure vlan vlan_name router-discovery {ipv6} hop-limit
```

Description

Unconfigures the current hop limit value sent in router discovery advertisements on the VLAN.

Syntax Description

<i>vlan_name</i>	Specifies an IPv6 configured VLAN.
------------------	------------------------------------

Default

N/A.

Usage Guidelines

This command sets the hop-limit parameter to the default value of 64.

Example

The following example unconfigures the current hop limit for the VLAN "top_floor":

```
unconfigure vlan top_floor router-discovery hop-limit
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

unconfigure vlan router-discovery link-mtu

```
unconfigure vlan vlan_name router-discovery {ipv6} link-mtu
```

Description

Unconfigures the link MTU value sent in router discovery advertisements on the VLAN.

Syntax Description

<i>vlan_name</i>	Specifies an IPv6 configured VLAN.
------------------	------------------------------------

Default

N/A.

Usage Guidelines

This command sets the **link-mtu** parameter to the default value of 0.

Example

The following example unconfigures the link MTU for the VLAN "top_floor":

```
unconfigure vlan top_floor router-discovery link-mtu
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

unconfigure vlan router-discovery managed-config-flag

```
unconfigure vlan vlan_name router-discovery {ipv6} managed-config-flag
```

Description

Unconfigures the managed address configuration flag value sent in router discovery advertisements on the VLAN.

Syntax Description

<i>vlan_name</i>	Specifies an IPv6 configured VLAN.
------------------	------------------------------------

Default

N/A.

Usage Guidelines

This command sets the managed-config-flag parameter to the default value off.

Example

The following example unconfigures the managed address configuration flag for the VLAN "top_floor":

```
unconfigure vlan top_floor router-discovery managed-config-flag
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

unconfigure vlan router-discovery max-interval

```
unconfigure vlan vlan_name router-discovery {ipv6} max-interval
```

Description

Unconfigures the maximum time between unsolicited router discovery advertisements on the [VLAN](#).

Syntax Description

<code>vlan_name</code>	Specifies an IPv6 configured VLAN.
------------------------	------------------------------------

Default

N/A.

Usage Guidelines

This command sets the max-interval parameter to the default value of 600 seconds.

Example

The following example unconfigures the max-interval for the VLAN "top_floor":

```
unconfigure vlan top_floor router-discovery max-interval
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

unconfigure vlan router-discovery min-interval

```
unconfigure vlan vlan_name router-discovery {ipv6} min-interval
```

Description

Unconfigures the minimum time between unsolicited router discovery advertisements on the [VLAN](#).

Syntax Description

<code>vlan_name</code>	Specifies an IPv6 configured VLAN.
------------------------	------------------------------------

Default

N/A.

Usage Guidelines

This command sets the min-interval parameter to the default value of (max-interval × .33 seconds).

Example

The following example unconfigures the min-interval for the VLAN "top_floor":

```
unconfigure vlan top_floor router-discovery min-interval
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

unconfigure vlan router-discovery other-config-flag

```
unconfigure vlan vlan_name router-discovery {ipv6} other-config-flag
```

Description

Unconfigures the other stateful configuration flag value sent in router discovery advertisements on the VLAN.

Syntax Description

<i>vlan_name</i>	Specifies an IPv6 configured VLAN.
------------------	------------------------------------

Default

N/A.

Usage Guidelines

This command sets the other-config-flag parameter to the default value off.

Example

The following example unconfigures the other stateful configuration flag for the VLAN "top_floor":

```
unconfigure vlan top_floor router-discovery other-config-flag
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

unconfigure vlan router-discovery reachable-time

```
unconfigure vlan vlan_name router-discovery {ipv6} reachable-time
```

Description

Unconfigures the reachable time value in router discovery advertisements on the VLAN.

Syntax Description

<i>vlan_name</i>	Specifies an IPv6 configured VLAN.
------------------	------------------------------------

Default

N/A.

Usage Guidelines

This command sets the reachable-time parameter to the default value of 30,000 milliseconds.

Example

The following example unconfigures the reachable time for the VLAN top_floor:

```
unconfigure vlan top_floor router-discovery reachable-time
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

unconfigure vlan router-discovery retransmit-time

```
unconfigure vlan vlan_name router-discovery {ipv6} retransmit-time
```

Description

Unconfigures the retransmit time value in router discovery advertisements on the VLAN.

Syntax Description

<code>vlan_name</code>	Specifies an IPv6 configured VLAN.
------------------------	------------------------------------

Default

N/A.

Usage Guidelines

This command sets the retransmit-time parameter to the default value of 1000 milliseconds.

Example

The following example unconfigures the retransmit time for the VLAN "top_floor":

```
unconfigure vlan top_floor router-discovery retransmit-time
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature in the [Switch Engine 32.3 Feature License Requirements](#) document.

unconfigure vlan subvlan-address-range

```
unconfigure vlan vlan_name subvlan-address-range
```

Description

Unconfigures subVLAN address ranges on each subVLAN to prohibit the entry of IP addresses from hosts outside of the configured range.

Syntax Description

<code>vlan_name</code>	Specifies a subVLAN name.
------------------------	---------------------------

Default

N/A.

Usage Guidelines

This command removes a subVLAN address range. There is no error checking to prevent the configuration of overlapping subVLAN address ranges between multiple subVLANs. Doing so can result in unexpected behavior of ARP within the superVLAN and associated subVLANs.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure vlan udp-profile

```
unconfigure vlan vlan_name udp-profile
```

Description

Removes any UDP forwarding profile from a VLAN.

Syntax Description

<i>vlan_name</i>	Specifies a VLAN name.
------------------	------------------------

Default

No UDP profiles are associated with the VLAN.

Usage Guidelines

None.

Example

The following example removes any UDP forwarding profile from the VLAN "to-sales":

```
unconfigure vlan to-sales udp-profile
```

History

This command was first available in ExtremeXOS 11.2.

Platform Availability

This command is available on all platforms that use the Edge, Advanced Edge, or Core license. For information on the licenses available for each platform, see the [Switch Engine 32.3 Feature License Requirements](#) document.

unconfigure vman ethertype

```
unconfigure vman ethertype {secondary}
```

Description

Restores the default primary VMAN ethertype value of 0x88A8 or deletes the secondary ethertype value.

Syntax Description

secondary	Deletes the secondary ethertype value.
------------------	--

Default

If the secondary option is not specified, it restores the default primary VMAN ethertype value of 0x88a8.

Usage Guidelines

When you enter this command without the secondary option, the primary VMAN ethertype returns to the default value of 0x88A8. If you specify the secondary option, the secondary VMAN ethertype value is deleted (no value is assigned).



Note

Before unconfiguring the secondary VMAN ethertype, any secondary VMAN port must be changed to the primary VMAN ethertype; otherwise the command fails.

Example

The following example restores the primary VMAN ethertype to the default value:

```
unconfigure vman ethertype
```

The following example deletes the secondary VMAN ethertype:

```
unconfigure vman ethertype secondary
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure vm-tracking local-vm

```
unconfigure vm-tracking local-vm mac-address mac [name | ip-address |  
vpp | vlan-tag]
```

Description

Unconfigures the parameters associated with a local VM database entry to be used for VM MAC local authentication.

Syntax Description

<i>mac</i>	Specifies the MAC address for the local VM database entry you want to unconfigure.
name	Removes the name configured for the VM database entry.
ip-address	Removes the IP address configured for the VM database entry.
vpp	Removes the VPP configured for the VM database entry.
vlan-tag	Removes the <u>VLAN</u> tag configured for the VM database entry.

Default

N/A.

Usage Guidelines

None.

Example

The following command removes the IP address configuration for the VM entry specified by the MAC address:

```
# unconfigure vm-tracking local-vm mac-address 00:E0:2B:12:34:56 ip-address
```

History

This command was first available in ExtremeXOS 12.5.

The ingress-vpp and egress-vpp options were replaced with the vpp option in ExtremeXOS 12.6.

The VLAN-tag option was added in ExtremeXOS 15.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure vm-tracking nms

```
unconfigure vm-tracking nms {server [primary | secondary]}
```

Description

Removes the configuration for one or both NMS servers.

Syntax Description

primary secondary	Specifies the whether you are unconfiguring the primary or secondary NMS.
-----------------------------------	---

Default

N/A.

Usage Guidelines

If you do not specify primary or secondary, this command removes the configuration for both NMS servers.

Example

The following command removes the configuration for the secondary NMS server:

```
# unconfigure vm-tracking nms server secondary
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure vm-tracking repository

```
unconfigure vm-tracking repository {primary | secondary}
```

Description

Removes the configuration for FTP file synchronization for NVPP and VMMAP files.

Syntax Description

primary secondary	Specifies the whether you are unconfiguring the primary or secondary FTP server.
-----------------------------------	--

Default

If you do not specify primary or secondary, the default action is to remove both the primary and secondary FTP server configurations.

Usage Guidelines

None.

Example

The following command removes the configuration for the primary and secondary FTP servers:

```
# unconfigure vm-tracking repository
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure vm-tracking vpp vlan-tag

```
unconfigure vm-tracking vpp vpp_name vlan-tag
```

Description

Unconfigures the VLAN tag of VPP.

Syntax Description

<i>vpp_name</i>	Specifies a name of the VPP.
-----------------	------------------------------

Default

N/A.

Usage Guidelines

Use this command to unconfigure the VLAN tag of VPP.

Example

Example output not yet available and will be provided in a future release.

History

This command was first available in ExtremeXOS 15.3.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure vm-tracking vpp

```
unconfigure vm-tracking vpp vpp_name
```

Description

Removes the association of a policy or ACL rule to an LVPP.

Syntax Description

<i>vpp_name</i>	Specifies the name of an existing LVPP.
-----------------	---

Default

N/A.

Usage Guidelines

None.

Example

The following command removes the configuration of LVPP vpp1:

```
# unconfigure vm-tracking vpp vpp1
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure vpex

```
unconfigure vpex ports port_list slot
```

Description

Allows you to disassociate a bridge port extender (BPE) from a slot number assignment.

Syntax Description

vpex	Specifies Virtual Port Extender (VPEX).
ports	Specifies switch ports attached to BPEs.
<i>port_list</i>	Specifies switch ports attached to BPEs. Must be in the format <i>slot:port</i> . Note: If the switch port is a LAG, the port specified must be the master port.
slot	Specifies the BPE slot assignment.

Default

N/A.

Usage Guidelines

The behavior of this command is similar to removing slots within a chassis. After executing this command, the BPE no longer occupies the slot number assignment.

Example

The following example removes the slot number assignment from a BPE attached to switch port 1:23:

```
# unconfigure vpex ports 1:23 slot
```

History

This command was first available in ExtremeXOS 22.5.

Platform Availability

This command is available on ExtremeSwitching 5420 and 5520 series switches.

unconfigure vpex mlag-id peer

```
unconfigure vpex mlag-id mlag_id slot
```

Description

In an Extended Edge Switching topology, removes the bridge port extender (BPE) slot assignment applied to an identifier on the specified MLAG peer when the port connected to the BPE is physically connected to the MLAG peer switch.

Syntax Description

vpex	Specifies Virtual Port Extender (VPEX).
mlag-id	Specifies a unique MLAG identifier of the MLAG port attached to the bridge port extender (BPE).
<i>mlag_id</i>	Specifies the MLAG identifier value of the MLAG port attached to the BPE. Range is 1-65,000.
slot	Specifies unconfiguring the slot identifier for the attached BPE.

Default

N/A.

Usage Guidelines

The configure **vpex mlag-id mlag_id peer peer_name slot slot_num** command allows the MLAG peer that does not have a port in the specified MLAG to declare a Extended Edge Switching slot on that MLAG. This command enables you to unconfigure a previously declared slot on that MLAG.

Example

The following example unconfigures the slot previously declared on MLAG "11":

```
# unconfigure vpex mlag-id 11 slot
```

History

This command was first available in ExtremeXOS 22.7.

Platform Availability

This command is available on ExtremeSwitching 5420 and 5520 series switches.

unconfigure vpls dot1q ethertype

unconfigure vpls *vpls_name* **dot1q ethertype**



Note

This command has been replaced with the following command: `unconfigure l2vpn [vpls vpls_name | vpwsvpws_name] dot1q ethertype`.

This command is still supported for backward compatibility, but it will be removed from a future release, so we recommend that you start using the new command.

Description

Unconfigures the ethertype setting for the VPLS specified by *vpls_name*.

Syntax Description

<i>vpls_name</i>	Identifies the VPLS within the switch (character string)
------------------	--

Default

N/A

Usage Guidelines

This command unconfigures the ethertype setting for the VPLS specified by *vpls_name*. The setting is changed back to the value displayed in the `show dot1q` command.

Example

The following command changes the ethertype setting for the specified VPLS to the value displayed in the `show dot1q` command:

```
unconfigure vpls my_vpls dot1q ethertype
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available only on the platforms that support *MPLS* as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

unconfigure vpls snmp-vpn-identifier

```
unconfigure vpls vpls_name snmp-vpn-identifier
```

Description

Removes an [SNMP](#) VPN identifier for traps from the specified VLPLS.

Syntax Description

<i>vpls_name</i>	Specifies the VPLS for which you are removing the identification string.
------------------	--

Default

N/A.

Usage Guidelines

None.

Example

The following command removes the identifier for SNMP VPN traps on VPLS vpls1:

```
unconfigure vpls vpls1 snmp-vpn-identifier
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available only on the platforms that support [MPLS](#) as described in the [Switch Engine 32.3 Feature License Requirements](#) document.

unconfigure vr description

```
unconfigure vr name description
```

Description

Removes a description for the specified VR or VRF.

Syntax Description

<i>name</i>	Specifies the name of a user VR or a VRF.
-------------	---

Default

No description.

Usage Guidelines

None.

Example

The following example removes a description for the VRF "corporate":

```
unconfigure vr corporate description
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure vr rd

```
unconfigure vr vrf_name rd
```

Description

This command removes the configuration for a VPN VRF RD.

Syntax Description

<i>vrf_name</i>	Specifies the name of a VPN VRF.
-----------------	----------------------------------

Default

N/A.

Usage Guidelines

None.

Example

The following examples unconfigure RDs using the two of the supported formats:

```
unconfigure vr corporate-extreme rd
unconfigure vr corporate-guest rd
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure vr vpn-id

```
unconfigure vr vrf_name vpn-id
```

Description

This command removes the configuration for a globally unique identifier for a VPN VRF.

Syntax Description

<i>vrf_name</i>	Specifies the name of a VPN VRF.
-----------------	----------------------------------

Default

N/A.

Usage Guidelines

None.

Example

The following example removes VPN ID ac:9f3c8 from the VRF "corporate-extreme":

```
unconfigure vr corporate-extreme vpn-id
```

History

This command was first available in ExtremeXOS 12.5.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

unconfigure xml-notification

```
unconfigure xml-notification
```

Description

Unconfigures the XML notification client.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to unconfigure the XML client process including the associated log target configuration.

Example

The following command unconfigures the xml-notification client:

```
unconfigure xml-notification
```

History

This command was first available in ExtremeXOS 12.4.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

uninstall image

```
uninstall image fname partition {msm slotid} {reboot}
```

On a SummitStack, use:

```
uninstall image fname partition {slot slot number} {reboot}
```

Description

Uninstalls an ExtremeXOS software package. Also uninstalls a VOSS image when staged for installation on the next reboot.

Syntax Description

<i>fname</i>	Specifies the software package to uninstall.
<i>partition</i>	Specifies which partition the package was installed to: primary or secondary. Select primary to remove it from the primary partition and secondary to remove it from the secondary partition.
<i>slotid</i>	On a SummitStack, the slot number specifies the node on which the BootROM image should be uninstalled.
reboot	Reboots the switch after the package is uninstalled.

Default

N/A.

Usage Guidelines

Use this command to uninstall a software package previously installed on the switch.

Use this command to uninstall a VOSS image when staged for installation on the next reboot.

```
* (NOS Change to VOSS after Reboot) switch-model-EXOS.3 # uninstall image
Uninstalling VOSS image...
Image uninstalled successfully
* (Beta) switch-model-EXOS.4 #
```

If a VOSS image is not staged, the same command would yield the following error:

```
* (Beta) switch-model-EXOS.4 # uninstall image
Uninstalling VOSS image...
Error: No other Network Operating System image has been staged for auto-install
* (Beta) switch-model-EXOS.4 #
```

When you uninstall a software package, the switch prompts you to save your changes to your currently active configuration file:

```
Uninstallation of the EXOS module Do you want to save configuration changes to
primary.cfg? (y or n)
```

Enter **y** to save the changes to your configuration file. Enter **n** to not save the changes to your configuration file.

Local Filename Character Restrictions

This section provides information about the characters supported by the switch for local and remote filenames.

When specifying a local filename, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-)
- Underscore (_)

When naming a local file, remember the requirements previously described.

SummitStack Only

You can issue this command only from a Master node.

Example

```
# uninstall image *switch-ntp.xmod secondary
```

History

This command was first available in ExtremeXOS 11.0.

The slot parameter was added to support SummitStack in ExtremeXOS 12.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

uninstall license file

```
uninstall license file filename [{revoke revocation_file} | withhold ]
    {slot slot}
```

Description

Uninstalls a license key file on ExtremeSwitching 5320, 5420 and 5520 series switches.

Syntax Description

file	Specifies providing the license key file.
<i>filename</i>	Specifies the name of the license key file.
slot	For stacks and Extended Edge Switching, specifies uninstalling the license key file from a slot.
<i>slot</i>	For stacks and Extended Edge Switching, specifies the slot to uninstall the license key file from.
revoke	Specifies generating a revocation certificate to release the license entitlement and invalidate the license key.
<i>revocation_file</i>	Specifies the file name (with <code>.rvk</code> extension) to store the revocation certificate in.
withhold	Specifies retaining the license entitlement for subsequent re-installation with the original license file. If the original license file is not available, use the revoke option.

Default

N/A.

Usage Guidelines

This command disables and uninstalls the feature license contained in the RMS license file. This command accepts *filename* as a `.lic` file that contains one or more licenses to be uninstalled. This is the license file that was used to install license. Only files with `.lic` extension are accepted.

The **revoke** option invalidates the feature license and generates a revocation certificate, which is the first step to release the license entitlement back to the license entitlement manager (LEM), whereas the **withhold** option retains the license entitlement with the switch, and thus, the original license can be used to reinstall the feature license. If the specified file name (for *revocation_file*) does not have an `.rvk` extension, it is automatically appended.

If you use this command to disable and remove the license, this removes the configuration capability of related features. However, the functionality continues to work until the next reboot.

To install a license, use the command `install license file filename {slot slot}`.

To uninstall a specific license product, use the command `uninstall license product product_name [revoke revocation_file | withhold] {slot slot}`

To view the licenses installed on your switch, use the `show licenses {[slot slot |all]} {detail}` command.

Example

The following example uninstalls the license with filename `mylicense.lic`, so that it can be re-installed on the switch at a future date:

```
# uninstall license file mylicense.lic withhold
```

History

This command was first available in ExtremeXOS 31.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

uninstall license product

```
uninstall license product product_name [revoke revocation_file |  
    withhold] {slot slot}
```

Description

Disables and uninstalls specific license products for ExtremeSwitching all Universal switches.



Note

licMgr does not recognize PRD-5000-MACSEC feature license on ExtremeSwitching 5420 series switches.

Syntax Description

product	Specifies uninstalling a particular license product.
<i>product_name</i>	Selects the specific license product to uninstall.
slot	For stacks and Extended Edge Switching, specifies uninstalling the license feature from a slot.
<i>slot</i>	For stacks and Extended Edge Switching, specifies the slot to uninstall the license feature from.
revoke	Specifies generating a revocation certificate to release the license entitlement and invalidate the license key.
<i>revocation_file</i>	Specifies the file name (with .rvk extension) to store the revocation certificate in.
withhold	Specifies retaining the license entitlement for subsequent re-installation with the original license file. If the original license file is not available, use the revoke option.

Default

N/A.

Usage Guidelines

This command uninstalls a specific feature license product. This command effectively disables the specified feature.

This command accepts *product_name* as the name of the license product to be uninstalled. To view the license products installed on your switch, use the `show licenses {[slot slot |all]} {detail}` command with the **detail** option.

The **revoke** option invalidates the feature license and generates a revocation certificate, which is the first step to release the license entitlement back to the license entitlement manager (LEM), whereas the **withhold** option retains the license entitlement with the switch, and thus, the original license can be used to reinstall the feature license. If the specified file name (for *revocation_file*) does not have an .rvk extension, it is automatically appended.

To install a license, use the following command `install license file filename {slot slot}`.

Example

The following example removes the MACsec (PRD-5000-MACSEC) license product from a switch: and generate a revocation certificate (macsec_lic.rvk) to release the license entitlement and invalidate the license key:

```
# uninstall license product PRD-5000-MACSEC revoke macsec_lic.rvk
```

History

This command was first available in ExtremeXOS 31.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

upload configuration

```
upload configuration [hostname | ipaddress] filename {vr vr-name}
                {block-size block_size}
```

Description

Uploads the current configuration in ASCII format to a TFTP server on your network.

Syntax Description

<i>hostname</i>	Specifies the hostname of the TFTP server where you want to download the configuration file. You must have DNS enabled.
<i>ipaddress</i>	Specifies the IP address of the TFTP server where you want to download the configuration file.
<i>filename</i>	Specifies a user-defined name for the configuration file. You must use the .xsf file extension when naming an ASCII-formatted configuration file.
<i>vr-name</i>	Specifies the name of the virtual router. By default the switch uses <i>VR-Mgmt</i> for this command. NOTE: User-created VRs are supported only on the platforms listed for this feature in the Switch Engine 32.3 Feature License Requirements document.
<i>block_size</i>	Specifies the data block size, excluding TFTP header. Data block size ranges from 24-65000 bytes.

Default

Uploads the current configuration in ASCII format immediately to a TFTP server.

The default value of **block-size** is 1400 bytes.

Usage Guidelines

Specify the `ipaddress` or `hostname` parameters to upload the current, active configuration file from the switch to a TFTP server on the network. Use of the `hostname` parameter requires that DNS be enabled.

The uploaded ASCII file retains the 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.

This command is not applicable to XML-based configurations. Those files use the `.cfg` file extension.

If you want to view your configuration in ASCII format, use the `.xsf` file extension (known as the XOS script file) when you save the configuration file on the switch. This saves the XML-based configuration in an ASCII format readable by a text editor.

If you successfully upload the active configuration to the network TFTP server, the switch displays a message similar to the following:

```
Uploading meg_upload_config1.xsf to 10.10.10.10 ... done!
```

If the switch displays a timeout error message similar to the following:

```
failed! Error: timeout
```

Make sure you entered the correct host name or IP address of the TFTP server

If the switch displays an unreachable network error similar to the following:

```
failed! Error: Network is unreachable
```

Make sure you entered the correct virtual router. By default the switch uses VR-Mgmt for this command.

Summary of Steps

The following summary only describes the CLI involved to transfer the configuration and load it on the switch; it is assumed that you know how to modify the configuration file with a text editor. As previously described, to use these commands, use the `.xsf` file extension. These steps are not applicable to configurations that use the `.cfg` file extension.

To work with an ASCII-formatted configuration file, complete the following tasks:

- Upload the configuration to a network TFTP server using the following command:

```
upload configuration [hostname | ipaddress] filename {vr vr-name}  
{block-size block_size}
```

After the configuration file is on the TFTP server, use a text editor to the desired changes.

- Download the configuration from the TFTP server to the switch using one of the following commands:

```
tftp [host-name | ip-address] -g -r remote-file
```

```
tftp get [host-name | ip-address] remote-file
```

- Verify the configuration file is on the switch using the following command:

```
ls
```

- Load and restore the new configuration file on the switch using the following command:

```
load script filename {arg1} {arg2} ... {arg9}
```

- Save the configuration to the configuration database so the switch can reapply the configuration after switch reboot using the following command:

```
save configuration {primary | secondary | existing-config | new-config}
```

When you save the configuration file, the switch automatically adds the .cfg file extension to the filename. This saves the ASCII configuration as an XML-based configuration file.

The following describes the steps in more detail.

Uploading the ASCII Configuration File To a TFTP Server

To upload the current switch configuration as an ASCII-based file to the TFTP server, use the `upload configuration` command and save the configuration with the .xsf file extension.

For example, to transfer the current switch configuration as an ASCII-based file named meg_upload_config1.xsf to the TFTP server with an IP address of 10.10.10.10, do the following:

```
upload configuration 10.10.10.10 meg_upload_config1.xsf
```

If you successfully upload the configuration to the TFTP server, the switch displays a message similar to the following:

```
Uploading meg_upload_config1.xsf to 10.10.10.10 ... done!
```

Downloading the ASCII Configuration File to the Switch

To download the configuration from the TFTP server to the switch, use the `tftp` command. For example, to retrieve the configuration file named meg-upload_config1.xsf from a TFTP server with an IP address of 10.10.10.10, you can use one of the following commands:

```
tftp 10.10.10.10 -g -r meg_upload_config1.xsf  
tftp get 10.10.10.10 meg_upload_config1.xsf
```

If you successfully download the configuration to the switch, the switch displays a message similar to the following:

```
Downloading meg_upload_config1.xsf to switch... done!
```

Verifying that the ASCII Configuration File is on the Switch

To confirm that the ASCII configuration file is on the switch, use the `ls` command. The file with an .xsf extension is the ASCII configuration.

The following sample output contains an ASCII configuration file:

```
-rw-r--r-- 1 root 0 98362 Nov 2 13:53 Nov022005.cfg
-rw-r--r-- 1 root 0 117136 Dec 12 12:56 epicenter.cfg
-rw-r--r-- 1 root 0 68 Oct 26 11:17 mcastgroup.pol
-rw-r--r-- 1 root 0 21203 Dec 13 15:40 meg_upload_config1.xsf
-rw-r--r-- 1 root 0 119521 Dec 6 14:35 primary.cfg
-rw-r--r-- 1 root 0 96931 Nov 11 11:01 primary_11_11_05.cfg
-rw-r--r-- 1 root 0 92692 Jul 19 16:42 secondary.cfg
```

Loading the ASCII Configuration File

After downloading the configuration file, you must load the new configuration on the switch. To load and restore the ASCII configuration file, use the `load script filename {arg1} {arg2} ... {arg9}` command. After issuing this command, the ASCII configuration quickly scrolls across the screen.

The following is an example of the type of information displayed when loading the ASCII configuration file:

```
script.meg_upload_config1.xsf.389 # enable snmp access
script.meg_upload_config1.xsf.390 # enable snmp traps
script.meg_upload_config1.xsf.391 # configure mstp region purple
script.meg_upload_config1.xsf.392 # configure mstp revision 3
script.meg_upload_config1.xsf.393 # configure mstp format 0
script.meg_upload_config1.xsf.394 # create stpd s0
```

Instead of entering each command individually, the script runs and loads the CLI on the switch.

Saving the Configuration

After you load the configuration, save it to the configuration database for use by the switch. This allows the switch to reapply the configuration after a switch reboot. To save the configuration, use the `save configuration {primary | secondary | existing-config | new-config}` command.

When you save the configuration file, the switch automatically adds the .cfg file extension to the filename. This saves the ASCII configuration as an XML-based configuration file.

You can use any name for the configuration. For example, after loading the file meg_upload_config1.xsf, you need to save it to the switch. To save the configuration as configuration1.cfg, do the following:

```
save configuration configuration1
```

Host Name and Remote IP Address Character Restrictions

This section provides information about the characters supported by the switch for host names and remote IP addresses.

When specifying a host name or remote IP address, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)

- Period (.)
- Dash (-) Permitted only for host names
- Underscore (_) Permitted only for host names
- Colon (:)

When naming or configuring an IP address for your network server, remember the requirements listed above.

Remote Filename Character Restrictions

This section provides information about the characters supported by the switch for remote filenames.

When specifying a remote filename, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-)
- Underscore (_)
- Slash (/) Permitted only for remote files

When naming a remote file, remember the requirements previously described.

Example

The following command uploads the current switch configuration as an ASCII-based file named configbackup.xsf to the TFTP server with an IP address of 10.10.10.10:

```
upload configuration 10.10.10.10 configbackup.xsf
```

If you successfully upload the configuration to the TFTP server, the switch displays a message similar to the following:

```
Uploading configbackup.xsf to 10.10.10.10 ... done!
```

History

This command was first available in ExtremeXOS 11.4.

Block size support was added in ExtremeXOS 15.7.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

upload debug

```
upload debug [hostname | ipaddress] [{vr} vr_name] {block-size  
block_size}
```

Description

Uploads debug information files to a tftp server.

Syntax Description

<i>hostname</i>	Specifies the host name of the TFTP server to which the debug files will be uploaded to.
<i>ipaddress</i>	Specifies the IP address of the TFTP server to which the debug files will be uploaded to.
<i>vr_name</i>	Specifies the name of the virtual router.
<i>block_size</i>	Specifies the data block size, excluding TFTP header. Data block size ranges from 24-65000 bytes.

Default

By default, the virtual router VR-Mgmt will be used.

If you do not specify block size, the default value is 1400 bytes.

Usage Guidelines



Note

Use this command only under the guidance of Extreme Networks Technical Support personnel to troubleshoot the switch.

Use this command to copy, upload debug information (for example, core, trace, show tech-support, configuration, and policy files) to the specified TFTP server.

Progress messages are displayed that indicate the file being copied and when the copying is finished. Depending on your platform, the switch displays a message similar to the following:

```
The following files on have been uploaded: Tarball Name: TechPubsLab_C_09271428.tgz ./
primary.cfg
```

You can also use this command in conjunction with the `show tech-support` command. Prior to uploading debug information files, the switch prompts you with the following message to run the `show tech` command with the logto file option:

```
Do you want to run show tech logto file first? (y/n)
```

Enter y to run the `show tech-support` command before uploading debug information. If you enter y, the `show_tech.log.tgz` file is included during the upload. Enter n to upload debug information without running the `show tech` command.

After you upload the debug information, you should see a compressed TAR file, which contains the debug information.

The TAR file naming convention is:

```
<SysName>_<{<slot#>|A|B}I|C>_<Current Time>.tgz
```

where <Current Time> = mmddhhmm (month (01-12), day (01-31), hour (0-24), minute (00-59)).

Example

The following command uploads debug files to a network TFTP server:

```
upload debug 10.10.10.10
```

History

This command was first available in ExtremeXOS 11.6.

Block size support was added in ExtremeXOS 15.7.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

upload dhcp-bindings

```
upload dhcp-bindings
```

Description

Upload the DHCP bindings immediately on demand.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This commands enables the functionality to allow you to upload DHCP bindings on demand.

History

This command was first available in ExtremeXOS 12.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

upload log

```
upload log [ ipaddress | hostname ] {vr vr_name} {block-size block_size}
      filename {messages [memory-buffer | nvram] } {severity severity
      {only}} {match regex} {chronological}
```

Description

Uploads the current log messages to a TFTP server.

Syntax Description

<i>ipaddress</i>	Specifies the ipaddress of the TFTP server.
<i>hostname</i>	Host name of server to send log information to.
<i>vr_name</i>	Specifies the virtual router that can reach the TFTP server. Note: User-created VRs are supported only on the platforms listed for this feature in the Switch Engine 32.3 Feature License Requirements document.
<i>block_size</i>	Specifies the data block size, excluding TFTP header. Data block size ranges from 24-65000 bytes.
<i>filename</i>	Specifies the file name for the log stored on the TFTP server.
messages	Specifies the location from which to display the log messages.
memory-buffer	Show messages stored in volatile memory.
nvr am	Show messages stored in NVRAM.
<i>severity</i>	Specifies the minimum severity level to display (if the keyword only is omitted).
only	Specifies that only the specified severity level is to be displayed.
<i>regex</i>	Specifies a regular expression. Only messages that match the regular expression will be displayed.
chronological	Specifies uploading log messages in ascending chronological order (oldest to newest).

Default

The following defaults apply:

- messages—memory buffer.
- severity—none (displays everything stored in the target).
- match—no restriction.
- chronological—if not specified, show messages in order from newest to oldest.

If you do not specify block size, the default value is 1400 bytes.

Usage Guidelines

This command is similar to the `show log` command, but instead of displaying the log contents on the command line, this command saves the log to a file on the TFTP server you specify. For more details on most of the options of this command, see the command `show log`.

Host Name and Remote IP Address Character Restrictions

This section provides information about the characters supported by the switch for host names and remote IP addresses.

When specifying a host name or remote IP address, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z).
- Numerals (0-9).
- Period (.).
- Dash (-) Permitted only for host names.
- Underscore (_) Permitted only for host names.
- Colon (:).

When naming or configuring an IP address for your network server, remember the requirements listed above.

Remote Filename Character Restrictions

This section provides information about the characters supported by the switch for remote filenames.

When specifying a remote filename, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z).
- Numerals (0-9).
- Period (.).
- Dash (-).
- Underscore (_).
- Slash (/).

When naming a local or remote file, remember the requirements listed above.

Example

The following command uploads messages with a critical severity to the filename switch4critical.log on TFTP server at 10.31.8.25:

```
upload log 10.31.8.25 switch4critical.log severity critical
```

The following command uploads messages with warning, error, or critical severity to the filename switch4warn.log on TFTP server at 10.31.8.25:

```
upload log 10.31.8.25 switch4warn.log severity warning
```

History

This command was first available in ExtremeXOS 10.1.

Block size support was added in ExtremeXOS 15.7.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

use configuration

```
use configuration file-name
```

Description

Configures the switch to use a previously saved configuration on the next reboot.

Syntax Description

<i>file-name</i>	Specifies an existing user-defined configuration file name (displays a list of available user-defined configuration files).
------------------	---

Default

N/A.

Usage Guidelines



Note

When the switch boots up after executing this command, examine the log to determine if a subsequent save and reboot of the switch or slot is necessary for the configuration to be correctly applied.

In rare cases, this command requires a second reboot for all configuration settings to take effect. This is because certain configuration settings are required for hardware initialization and are stored in the switch's non-volatile storage, as well as the configuration file. These settings include: configure forwarding internal-tables and configure ports partition. Therefore, when the configuration file selected by this command has different values for the above settings, a message is logged explaining that a save and reboot are required.

XML-based configuration files have a .cfg file extension. When you enter the name of the file in the CLI, the system automatically adds the .cfg file extension.

Do not use this command with ASCII-formatted configuration files. Those configuration files have an .xsf file extension. For more information about using and saving ASCII-formatted configuration files see the `upload configuration [hostname | ipaddress] filename {vr vr-name }` and the `load script filename {arg1} {arg2} ... {arg9}` commands.

There is no special significance to the primary and secondary configurations. They are just conveniences to specify the files `primary.cfg` and `secondary.cfg`.

When you configure the switch to use a previously saved configuration, the switch displays the following message:

The selected configuration file is now "<file-name>.cfg". By default, this file will be used for saving the configuration which will take effect after the next switch reboot.

You can create a new configuration file by saving your current switch configurations and using that file on the next reboot. For example, to create a new configuration named `test1` based on your current CLI session and switch configurations, use the following command:

```
save configuration test1
```

Tracking and Displaying Switch Configuration Files

To keep track of your configuration file names, use the `ls` command to display the files saved on your switch. Files with the `.cfg` extension are configuration files. In addition, you can see a list of available configuration files when you use the `use configuration` command.

The following is sample output from this command ("test" and "XOS1" are the names of the user-created and defined configurations):

```
exsh.1 # use configuration
primary      Primary configuration file
secondary    Secondary configuration file
<file-name>  Configuration file name
"test" "XOS1"
```

You can also use the `ls` command to display a list of the current configuration and policy files in the system.

Displaying the Active Configuration

To view the currently active, running configuration, use the `show switch` command.

Local Filename Character Restrictions

This section provides information about the characters supported by the switch for local filenames.

When specifying a local filename, the switch permits only the following characters:

- Alphabetical letters, upper case and lower case (A-Z, a-z)
- Numerals (0-9)
- Period (.)
- Dash (-)
- Underscore (_)

When naming a local file, remember the requirements listed above.

Example

The following command specifies that the next reboot should use the saved configuration file named XOS1.cfg:

```
use configuration XOS1
```

The following command specifies that the next reboot should use the configuration saved in the primary partition:

```
use configuration primary
```

History

This command was first available in ExtremeXOS 10.1.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

use image

```
use image {partition} partition {slot slotid}
```

Description

Configures the switch to use a saved image on the next reboot.

Syntax Description

<i>partition</i>	Specifies which image to use on the next reboot, the one stored on the primary partition, or the one stored on the secondary partition.
<i>slotid</i>	On a SummitStack, the slotid specifies the node on which the BootROM image is selected.

Default

The currently booted image.

Usage Guidelines

This command specifies which image to use on the next reboot. Two images can be stored, one on the primary partition, one on the secondary partition. To view your current (active) partition and the selected partition for the next reboot or installation, use the following command: `show switch`

Output from this command includes the selected and booted images and if they are in the primary or the secondary partition. Primary indicates the saved image in the primary partition; secondary indicates the saved image in the secondary partition.

SummitStack Only

You can issue this command only from a Master node. The image to use is stored in NVRAM on all target nodes.

If a slot number is not provided, the partition is selected on all nodes in the Active Topology.

Example

Using TFTP

The following command configures the switch to use the image stored in the primary partition on the next reboot:

```
use image partition primary
```

A message similar to the following is displayed:
To take effect of partition change please reboot the switch!

History

- This command was first available in ExtremeXOS 10.1.
- The **slot** parameter was added to support SummitStack in ExtremeXOS 12.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

virtual-router

```
virtual-router {vr-name}
```

Description

Changes the VR context.

Syntax Description

<i>vr-name</i>	Specifies the name of the VR.
----------------	-------------------------------

Default

N/A.

Usage Guidelines

Use this command to change the VR context for subsequent commands. When you issue the command, the prompt changes to reflect the VR domain. Configuration commands for Layer 3 routing protocols, creating VLANs, and deleting VLANs apply only to the current VR context.

Use this command with no name, or use the name "VR-Default" to return to the default configuration domain.

Under a VR configuration domain, any VR commands are applied only to that VR. The VR commands consist of all the BGP, OSPF, PIM, and RIP commands, and the commands listed in the following table.

Table 49: VR Commands

<code>[enable disable] ipforwarding</code>
<code>clear iparp</code>
<code>clear counters iparp⁴</code>
<code>configure iparp⁴</code>
<code>configure iparp [add delete] ⁴</code>
<code>[enable disable] iparp ⁴</code>
<code>show iparp ⁴</code>
<code>configure iproute [add delete] ⁴</code>
<code>show iproute ⁴</code>
<code>show ipstats ⁴</code>
<code>rtlookup</code>
<code>create [vlan vman] <i>vlan-name</i></code>
<code>[enable disable] igmp</code>
<code>[enable disable] igmp snooping ⁴</code>
<code>[enable disable] ipmcforwarding</code>
<code>show igmp</code>
<code>show igmp snooping</code>
<code>show igmp group</code>
<code>show igmp snooping cache</code>
<code>[enable disable] mld</code>
<code>[enable disable] mld snooping</code>
<code>show mld</code>
<code>show mld snooping</code>
<code>show mld group</code>

⁴ Other commands are available with these listed.

The VR context simplifies configuration because you do not have to specify the VR for each individual protocol configuration command. The current VR context is indicated in the CLI prompt.

For example, if you wish to configure OSPF for the user VR `vr-manufacturing`, you would change the VR context to that of `vr-manufacturing`. All the subsequent OSPF commands would apply to that VR, unless the context is changed again.

A VR is identified by a name (up to 32 characters long). The name must be unique among the *VLAN* and VR names on the switch. For backward compatibility, you cannot name a virtual router `VR-0`, `VR-1`, or `VR-2`. VR names are case insensitive.

When a new VR is created, by default, no ports are assigned, no VLAN interface is created, and no support for any routing protocols is added.

Example

The following example changes the VR context to `"vr-acme"`:

```
virtual-router vr-acme
```

History

This command was first available in ExtremeXOS 11.0.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

watch

```
watch command {difference} {count count} {interval interval}
```

Description

Runs a selected show command at a repeated interval.

Syntax Description

<i>command</i>	Specifies the command. Must be enclosed in double quotes.
difference	Shows the difference between successive iterations if numerical values have changed in the output.
count	Specifies the number of times to run the command.
<i>count</i>	Range is 1-4,294,967,295. The default is 3.
interval	Specifies the time between command iterations.
<i>interval</i>	The interval in seconds. Range is 1-4,294,967,295. The default is 5.

Default

The *count* default is 3.

The *interval* default is 5.

Usage Guidelines

This command runs the integrated Python script `watch.py` to run the specified show command periodically and display the output accordingly. In the script, CLI paging is disabled so that user input is not needed, and CLI refreshing is disabled so that the keyword **no-refresh** is not needed in auto-refreshing commands.

Example

The following example runs the show ports command for port 11 runs 15 times with an interval of 30 seconds:

```
# watch "show ports 11 statistics port-number no-refresh" count 15 interval 30
Port      Link      Tx Pkt    Tx Byte    Rx Pkt     Rx Byte    Rx Pkt     Rx
Pkt       Tx Pkt    Tx Pkt
Mcast     State     Count     Count      Count      Count      Bcast
Mcast     Bcast     Mcast
=====
=====

11         A         206845    30061110    205326     30939540    0
163381    0         163483

=====
=====

> in Port indicates Port Display Name truncated past 8 characters
> in Count indicates value exceeds column width. Use 'wide' option or '0' to
clear.

Link State: A-Active, R-Ready, NP-Port Not Present, L-Loopback

Port      Link      Tx Pkt    Tx Byte    Rx Pkt     Rx Byte    Rx Pkt     Rx
Pkt       Tx Pkt    Tx Pkt
Mcast     State     Count     Count      Count      Count      Bcast
Mcast     Bcast     Mcast
=====
=====

11         A         206845    30061110    205327     30939662    0
163382    0         163483

=====
=====

> in Port indicates Port Display Name truncated past 8 characters
```



```
> in Count indicates value exceeds column width. Use 'wide' option or '0' to
clear.

Link State: A-Active, R-Ready, NP-Port Not Present, L-Loopback

Port      Link      Tx Pkt      Tx Byte      Rx Pkt      Rx Byte      Rx Pkt      Rx
Pkt      Tx Pkt      Tx Pkt
Mcast      State      Count      Count      Count      Count      Bcast

=====
=====

11      A      206847      30061402      205328      30939832      0
163383      0      163485

=====
=====

> in Port indicates Port Display Name truncated past 8 characters

> in Count indicates value exceeds column width. Use 'wide' option or '0' to
clear.

Link State: A-Active, R-Ready, NP-Port Not Present, L-Loopback
```

History

This command was first available in ExtremeXOS 31.2.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.

WHILE ... DO

WHILE (**_expression**) DO



Note

This is a script command and operates only in scripts when scripting is enabled with the following command: [enable cli scripting {permanent}](#).

Description

Executes a code block while the specified condition is true.

Syntax Description

expression	Specifies the condition for which the statements should be executed while the condition is true.
<i>statements</i>	Set of statements to be executed while the condition is true.

Default

N/A.

Usage Guidelines

This command is usually followed by statements to be executed while the condition is true and the entire construct is terminated by an ENDWHILE command.

The `_expression` must be enclosed in parentheses.

Nesting is supported up to five levels. An error message is displayed if there is incorrect nesting of WHILE conditions. An error message is displayed if a user tries to execute more than five WHILE conditions.

Ctrl-C can be used to break out of a WHILE loop(s). Breaking out of any number of WHILE loops always clears all the WHILE loops .

The operators mentioned in [Using Operators](#) can be used in an `_expression` in a WHILE condition.

You can insert comments by using a number sign (#).

Example

This example creates 10 VLANs, named x1 to x10:

```
set var x 1

WHILE ($x <= 10) DO
    create vlan v$x
    set var x ($x + 1)
ENDWHILE
```

History

This command was first available in ExtremeXOS 11.6.

Platform Availability

This command is available on ExtremeSwitching 5320, 5420, 5520, and 5720 series switches.



Software Licensing

Extreme Networks software may contain software from third party sources that must be licensed under the specific license terms applicable to such software. Applicable copyright information is provided below.

Copyright (c) 1995-1998 by Cisco Systems, Inc.

Permission to use, copy, modify, and distribute this software for any purpose and without fee is hereby granted, provided that this copyright and permission notice appear on all copies of the software and supporting documentation, the name of Cisco Systems, Inc. not be used in advertising or publicity pertaining to distribution of the program without specific prior notice be given in supporting documentation that modification, permission, and copying and distribution is by permission of Cisco Systems, Inc.

Cisco Systems, Inc. makes no representations about the suitability of this software for any purpose. THIS SOFTWARE IS PROVIDED "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

MD5C.C - RSA Data Security, Inc., MD5 (Message-Digest algorithm 5) Message-Digest Algorithm

Copyright (C) 1991-2, RSA Data Security, Inc. Created 1991. All rights reserved.

License to copy and use this software is granted provided that it is identified as the "RSA Data Security, Inc. MD5 Message-Digest Algorithm" in all material mentioning or referencing this software or this function.

License is also granted to make and use derivative works provided that such works are identified as "derived from the RSA Data Security, Inc. MD5 Message-Digest Algorithm" in all material mentioning or referencing the derived work.

RSA Data Security, Inc. makes no representations concerning either the merchantability of this software or the suitability of this software for any particular purpose. It is provided "as is" without express or implied warranty of any kind.

These notices must be retained in any copies of any part of this documentation and/or software.

\$Id: md5c.c,v 1.2.4880.1 2005/06/24 01:47:07 lindak Exp \$ This code is the same as the code published by RSA Inc. It has been edited for clarity and style only.



Glossary

ACL

An Access Control List is a mechanism for filtering packets at the hardware level. Packets can be classified by characteristics such as the source or destination MAC, IP address, IP type, or QoS queue. Once classified, the packets can be forwarded, counted, queued, or dropped.

ad hoc mode

An 802.11 networking framework in which devices or stations communicate directly with each other, without the use of an AP.

ARP

Address Resolution Protocol is part of the TCP/IP suite used to dynamically associate a device's physical address (MAC address) with its logical address (IP address). The system broadcasts an ARP request, containing the IP address, and the device with that IP address sends back its MAC address so that traffic can be transmitted.

ATM

Asynchronous Transmission Mode is a start/stop transmission in which each character is preceded by a start signal and followed by one or more stop signals. A variable time interval can exist between characters. ATM is the preferred technology for the transfer of images.

BGP

Border Gateway Protocol is a router protocol in the IP suite designed to exchange network reachability information with BGP systems in other autonomous systems. You use a fully meshed configuration with BGP.

BGP provides routing updates that include a network number, a list of ASs that the routing information passed through, and a list of other path attributes. BGP works with cost metrics to choose the best available path; it sends updated router information only when one host has detected a change, and only the affected part of the routing table is sent.

BGP communicates within one AS using Interior BGP (IBGP) because BGP does not work well with IGP. Thus the routers inside the AS maintain two routing tables: one for the IGP and one for IBGP. BGP uses exterior BGP (EBGP) between different autonomous systems.

BSS

Basic Service Set is a wireless topology consisting of one access point connected to a wired network and a set of wireless devices. Also called an infrastructure network. See also [*IBSS \(Independent Basic Service Set\)*](#).

CDP

Cisco Discovery Protocol is a proprietary Data Link Layer protocol that shares information about other directly connected Cisco equipment, such as operating system versions and IP addresses.

Chalet

Chalet is a web-based user interface for setting up and viewing information about a switch, removing the need to enter common commands individually in the CLI.

CHAP

Challenge-Handshake Authentication Protocol is one of the two main authentication protocols used to verify a user's name and password for PPP Internet connections. CHAP is more secure because it performs a three-way handshake during the initial link establishment between the home and remote machines. It can also repeat the authentication anytime after the link has been established.

CLI

Command Line Interface. The CLI provides an environment to issue commands to monitor and manage switches and wireless appliances.

CoS

Class of Service specifies the service level for the classified traffic type.

Data Center Connect

DCC, formerly known as DCM (Data Center Manager), is a data center fabric management and automation tool that improves the efficiency of managing a large virtual and physical network. DCC provides an integrated view of the server, storage, and networking operations, removing the need to use multiple tools and management systems. DCC automates VM assignment, allocates appropriate network resources, and applies individual policies to various data objects in the switching fabric (reducing VM sprawl). Learn more about DCC at <http://www.extremenetworks.com/product/data-center-connect/>.

DHCP

Dynamic Host Configuration Protocol allows network administrators to centrally manage and automate the assignment of IP addresses on the corporate network. DHCP sends a new IP address when a computer is plugged into a different place in the network. The protocol supports static or dynamic IP addresses and can dynamically reconfigure networks in which there are more computers than there are available IP addresses.

DoS attack

Denial of Service attacks occur when a critical network or computing resource is overwhelmed so that legitimate requests for service cannot succeed. In its simplest form, a DoS attack is indistinguishable from normal heavy traffic. ExtremeXOS software has configurable parameters that allow you to defeat DoS attacks.

DSSS

Direct-Sequence Spread Spectrum is a transmission technology used in Local Area Wireless Network (LAWN) transmissions where a data signal at the sending station is combined with a higher data rate bit sequence, or chipping code, that divides the user data according to a spreading ratio. The chipping code is a redundant bit pattern for each bit that is transmitted, which increases the signal's resistance to interference. If one or more bits in the pattern are damaged during transmission, the original data

can be recovered due to the redundancy of the transmission. (Compare with *FHSS (Frequency-Hopping Spread Spectrum)*.)

EAP-TLS/EAP-TTLS

EAP-TLS Extensible Authentication Protocol - Transport Layer Security. A general protocol for authentication that also supports multiple authentication methods, such as token cards, Kerberos, one-time passwords, certificates, public key authentication and smart cards.

IEEE 802.1x specifies how EAP should be encapsulated in LAN frames.

In wireless communications using EAP, a user requests connection to a WLAN through an access point, which then requests the identity of the user and transmits that identity to an authentication server such as RADIUS. The server asks the access point for proof of identity, which the access point gets from the user and then sends back to the server to complete the authentication.

EAP-TLS provides for certificate-based and mutual authentication of the client and the network. It relies on client-side and server-side certificates to perform authentication and can be used to dynamically generate user-based and session-based WEP keys.

EAP-TTLS (Tunneled Transport Layer Security) is an extension of EAP-TLS to provide certificate-based, mutual authentication of the client and network through an encrypted tunnel, as well as to generate dynamic, per-user, per-session WEP keys. Unlike EAP-TLS, EAP-TTLS requires only server-side certificates.

(See also *PEAP (Protected Extensible Authentication Protocol)*.)

EAPS

Extreme Automatic Protection Switching is an Extreme Networks-proprietary version of the Ethernet Automatic Protection Switching protocol that prevents looping Layer 2 of the network. This feature is discussed in RFC 3619.

ECMP

Equal Cost Multi Paths is a routing algorithm that distributes network traffic across multiple high-bandwidth OSPF, BGP, IS-IS, and static routes to increase performance. The Extreme Networks implementation supports multiple equal cost paths between points and divides traffic evenly among the available paths.

EDP

Extreme Discovery Protocol is a protocol used to gather topology information about neighboring Extreme Networks switches.

ELRP

Extreme Loop Recovery Protocol is an Extreme Networks-proprietary protocol that allows you to detect Layer 2 loops.

EMS

The Event Management System is an Extreme Networks-proprietary system that saves, displays, and filters events, which are defined as any occurrences on a switch that generate a log message or require action.

ERPS

Ethernet Ring Protection Switching provides fast protection and recovery switching for Ethernet traffic in a ring topology. It also ensures that the Ethernet layer remains loop-free. It is defined in ITU/T G.8032.

ESRP

Extreme Standby Router Protocol is an Extreme Networks-proprietary protocol that provides redundant Layer 2 and routing services to users.

Extreme Defender for IoT

Extreme Defender for IoT provides unique in-line security for mission critical and/or vulnerable IoT devices. Placed between the IoT device and the network, the Defender for IoT solution helps secure and isolate IoT devices protecting them from internal and external hacking attempts, viruses, malware and ransomware, DDoS attacks, and more. Designed to be simple and flexible, Defender for IoT can be deployed over any network infrastructure to enable secure IoT management without significant network changes.

The solution is comprised of the Extreme Defender Application Software and the Defender Adapter (SA201) or AP3912i access point. ExtremeCloud Appliance is the supported platform for the Extreme Defender Application.

For more information, see <https://www.extremenetworks.com/product/extreme-defender-for-iot/>.

Extreme Management Center

Extreme Management Center (ExtremeCloud IQ - Site Engine), formerly Netsight™, is a web-based control interface that provides centralized visibility into your network. ExtremeCloud IQ - Site Engine reaches beyond ports, VLANs, and SSIDs and provides detailed control of individual users, applications, and protocols. When coupled with wireless and Identity & Access Management products, ExtremeCloud IQ - Site Engine becomes the central location for monitoring and managing all the components in the infrastructure. Learn more about ExtremeCloud IQ - Site Engine at <http://www.extremenetworks.com/product/management-center/>.

ExtremeAnalytics

ExtremeAnalytics™, formerly Purview™, is a network powered application analytics and optimization solution that captures and analyzes context-based application traffic to deliver meaningful intelligence about applications, users, locations, and devices. ExtremeAnalytics provides data to show how applications are being used. This can be used to better understand customer behavior on the network, identify the level of user engagement, and assure business application delivery to optimize the user experience. The software also provides visibility into network and application performance allowing IT to pinpoint and resolve performance issues in the infrastructure whether they are caused by the network, application, or server. Learn more about ExtremeAnalytics at <http://www.extremenetworks.com/product/extremeanalytics/>.

ExtremeCloud Appliance

The ExtremeCloud Appliance, the newest addition to the Smart OmniEdge portfolio, is a next generation orchestration application offering all the mobility services required for modern unified access deployments. The ExtremeCloud Appliance extends the simplified workflows of the ExtremeCloud public cloud application to on-prem/private cloud deployments.

The ExtremeCloud Appliance includes comprehensive critical network services for wireless and wired connectivity, wireless device secure onboarding, distributed and centralized data paths, role-based

access control through the Application Layer, integrated location services, and IoT device onboarding through a single platform.

Built on architecture with the latest technology, the embedded operating system supports application containers that enable future expansion of value added applications for the unified access edge. Learn more about ExtremeCloud Appliance at <https://www.extremenetworks.com/product/extremecloud-appliance/>.

ExtremeCloud

ExtremeCloud is a cloud-based network management Software as a Service (SaaS) tool. ExtremeCloud allows you to manage users, wired and wireless devices, and applications on corporate and guest networks. You can control the user experience with smarter edges – including managing QoS, call admission control, secure access policies, rate limiting, multicast, filtering, and traffic forwarding, all from an intuitive web interface. Learn more about ExtremeCloud at <http://www.extremenetworks.com/product/extremecloud/>.

ExtremeCloud™ IQ

ExtremeCloud™ IQ is an industry-leading and visionary approach to cloud-managed networking, built from the ground up to take full advantage of the Extreme Networks end-to-end networking solutions. ExtremeCloud IQ delivers unified, full-stack management of wireless access points, switches, and routers and enables onboarding, configuration, monitoring, troubleshooting, reporting, and more. Using innovative machine learning and artificial intelligence technologies, ExtremeCloud IQ analyzes and interprets millions of network and user data points, from the network edge to the data center, to power actionable business and IT insights, and deliver new levels of network automation and intelligence. Learn more about ExtremeCloud IQ at <https://www.extremenetworks.com/support/documentation/extremecloud-iq/>.

ExtremeControl

ExtremeControl, formerly Extreme Access Control™ (EAC), is a set of management software tools that use information gathered by a hardware engine to control policy to all devices on the network. The software allows you to automate and secure access for all devices on the network from a central dashboard, making it easier to roll out security and identity policies across the wired and wireless network. Learn more about ExtremeControl at <https://www.extremenetworks.com/product/extremecontrol/>.

ExtremeSwitching

ExtremeSwitching is the family of products comprising different switch types: **Modular** (X8 and 8000 series [formerly BlackDiamond] and S and K series switches); **Stackable** (X-series and A, B, C, and 7100 series switches); **Standalone** (SSA, X430, and D, 200, 800, and ISW series); and **Mobile Backhaul** (E4G). Learn more about ExtremeSwitching at <http://www.extremenetworks.com/products/switching-routing/>.

ExtremeWireless

ExtremeWireless products and solutions offer high-density WiFi access, connecting your organization with employees, partners, and customers everywhere they go. The family of wireless products and solutions includes APs, wireless appliances, and software. Learn more about ExtremeWireless at <http://www.extremenetworks.com/products/wireless/>.

ExtremeXOS

ExtremeXOS, a modular switch operating system, is designed from the ground up to meet the needs of large cloud and private data centers, service providers, converged enterprise edge networks, and everything in between. Based on a resilient architecture and protocols, ExtremeXOS supports network virtualization and standards-based SDN capabilities like VXLAN gateway and OpenStack Cloud orchestration. ExtremeXOS also supports comprehensive role-based policy.

FDB

The switch maintains a database of all MAC address received on all of its ports and uses this information to decide whether a frame should be forwarded or filtered. Each forwarding database (FDB) entry consists of the MAC address of the sending device, an identifier for the port on which the frame was received, and an identifier for the VLAN to which the device belongs. Frames destined for devices that are not currently in the FDB are flooded to all members of the VLAN. For some types of entries, you configure the time it takes for the specific entry to age out of the FDB.

FHSS

Frequency-Hopping Spread Spectrum is a transmission technology used in Local Area Wireless Network (LAWN) transmissions where the data signal is modulated with a narrowband carrier signal that 'hops' in a random but predictable sequence from frequency to frequency as a function of time over a wide band of frequencies. This technique reduces interference. If synchronized properly, a single logical channel is maintained. (Compare with [*DSSS \(Direct-Sequence Spread Spectrum\)*](#).)

IBSS

An IBSS is the 802.11 term for an ad hoc network. See [*ad hoc mode*](#).

ICMP

Internet Control Message Protocol is the part of the TCP/IP protocol that allows generation of error messages, test packets, and operating messages. For example, the ping command allows you to send ICMP echo messages to a remote IP device to test for connectivity. ICMP also supports traceroute, which identifies intermediate hops between a given source and destination.

IGMP

Hosts use Internet Group Management Protocol to inform local routers of their membership in multicast groups. Multicasting allows one computer on the Internet to send content to multiple other computers that have identified themselves as interested in receiving the originating computer's content. When all hosts leave a group, the router no longer forwards packets that arrive for the multicast group.

LAG

A Link Aggregation Group is the logical high-bandwidth link that results from grouping multiple network links in link aggregation (or load sharing). You can configure static LAGs or dynamic LAGs (using the LACP).

LLDP

Link Layer Discovery Protocol conforms to IEEE 802.1ab and is a neighbor discovery protocol. Each LLDP-enabled device transmits information to its neighbors, including chassis and port identification, system name and description, VLAN names, and other selected networking information. The protocol also specifies timing intervals in order to ensure current information is being transmitted and received.

MD5

Message-Digest algorithm is a hash function that is commonly used to generate a 128-bit hash value. It was designed by Ron Rivest in 1991. MD5 is officially defined in RFC 1321 - The MD5 Message-Digest Algorithm.

MIC

Message Integrity Check (or Code), also called 'Michael', is part of WPA and TKIP. The MIC is an additional 8-byte code inserted before the standard 4-byte ICV appended in by standard WEP to the 802.11 message. This greatly increases the difficulty in carrying out forgery attacks. Both integrity check mechanisms are calculated by the receiver and compared against the values sent by the sender in the frame. If the values match, there is assurance that the message has not been tampered with.

MLAG

The Multi-switch Link Aggregation Group feature allows users to combine ports on two switches to form a single logical connection to another network device. The other network device can be either a server or a switch that is separately configured with a regular LAG (or appropriate server port teaming) to form the port aggregation.

MPLS

Multiprotocol Label Switching speeds up network traffic. When forwarding packets, the Layer 2 (Switching) label is used to avoid complex destination lookups in the routing table. MPLS uses Label Switched Paths (LSPs) to establish the network path. The packet will be labeled so that service providers can decide the best way to keep traffic flowing. The Multiprotocol Label Switching Transport Profile (MPLS-TP) extensions to MPLS are designed to meet service provider requirements and are used as a network layer technology in transport networks. MPLS-TP gives service providers a reliable packet-based technology that is based on circuit-based transport networking. MPLS-TP is expected to be a low cost level 2 technology (if the limited profile is implemented in isolation) that will provide QoS, end-to-end OAM and protection switching.

MSDP

Multicast Source Discovery Protocol is used to connect multiple multicast routing domains. MSDP advertises multicast sources across Protocol Independent Multicast-Sparse Mode (PIM-SM) multicast domains or Rendezvous Points (RPs). In turn, these RPs run MSDP over TCP to discover multicast sources in other domains.

MSTI

Multiple Spanning Tree Instances control the topology inside an MSTP region. An MSTI is a spanning tree domain that operates within a region and is bounded by that region; and MSTI does not exchange BPDUs or send notifications to other regions. You can map multiple VLANs to an MSTI; however, each VLAN can belong to only one MSTI. You can configure up to 64 MSTIs in an MSTP region.

MSTP

Multiple Spanning Tree Protocol, based on IEEE 802.1Q-2003 (formerly known as IEEE 892.1s), allows you to bundle multiple VLANs into one STP topology, which also provides enhanced loop protection and better scaling. MSTP uses RSTP as the converging algorithm and is compatible with legacy STP protocols.

NetLogin

Network login provides extra security to the network by assigning addresses only to those users who are properly authenticated. You can use web-based, MAC-based, or IEEE 802.1X-based authentication with network login. The two modes of operation are campus mode and ISP mode.

netmask

A netmask is a string of 0s and 1s that mask, or screen out, the network part of an IP address, so that only the host computer part of the address remains. A frequently-used netmask is 255.255.255.0, used for a Class C subnet (one with up to 255 host computers). The ".0" in the netmask allows the specific host computer address to be visible.

OSPF

An interior gateway routing protocol for TCP/IP networks, Open Shortest Path First uses a link state routing algorithm that calculates routes for packets based on a number of factors, including least hops, speed of transmission lines, and congestion delays. You can also configure certain cost metrics for the algorithm. This protocol is more efficient and scalable than vector-distance routing protocols. OSPF features include least-cost routing, ECMP routing, and load balancing. Although OSPF requires CPU power and memory space, it results in smaller, less frequent router table updates throughout the network. This protocol is more efficient and scalable than vector-distance routing protocols.

OSPFv3

Open Shortest Path First version 3 is one of the routing protocols used with IPV6 and is similar to OSPF.

PEAP

Protected Extensible Authentication Protocol is an IETF draft standard to authenticate wireless LAN clients without requiring them to have certificates. In PEAP authentication, first the user authenticates the authentication server, then the authentication server authenticates the user. If the first phase is successful, the user is then authenticated over the SSL tunnel created in phase one using EAP-Generic Token Card (EAP-GTC) or Microsoft Challenged Handshake Protocol Version 2 (MSCHAP V2). (See also [EAP-TLS/EAP-TTLS](#).)

PoE

The Power over Ethernet standard (IEEE 802.3af) defines how power can be provided to network devices over existing Ethernet connections, eliminating the need for additional external power supplies.

QoS

Quality of Service is a technique that is used to manage network resources and guarantee a bandwidth relationship between individual applications or protocols. A communications network transports a multitude of applications and data, including high-quality video and delay-sensitive data such as real-time voice. Networks must provide secure, predictable, measurable, and sometimes guaranteed services. Achieving the required QoS becomes the secret to a successful end-to-end business solution.

RADIUS

RADIUS is a client/server protocol and software that enables remote access servers to communicate with a central server to authenticate dial-in users and authorize their access to the requested system or service. RADIUS allows a company to maintain user profiles in a central database that all remote servers can share. It provides better security, allowing a company to set up a policy that can be applied at a single administered network point. With RADIUS, you can track usage for billing and for keeping network statistics.

RIP

This IGP vector-distance routing protocol is part of the TCP/IP suite and maintains tables of all known destinations and the number of hops required to reach each. Using Routing Information Protocol, routers periodically exchange entire routing tables. RIP is suitable for use only as an IGP.

RIPng

Routing Information Protocol Next Generation is one of the routing protocols used with IPv6 and is similar to RIP.

SNMP

Simple Network Management Protocol is a standard that uses a common software agent to remotely monitor and set network configuration and runtime parameters. SNMP operates in a multivendor environment, and the agent uses MIBs, which define what information is available from any manageable network device. You can also set traps using SNMP, which send notifications of network events to the system log.

SNTP

Simple Network Time Protocol is used to synchronize the system clocks throughout the network. An extension of NTP, SNTP can usually operate with a single server and allows for IPv6 addressing.

SSL

Secure Socket Layer is a protocol for transmitting private documents using the Internet. SSL works by using a public key to encrypt data that is transferred over the SSL connection. SSL uses the public-and-private key encryption system, which includes the use of a digital certificate.

STP

Spanning Tree Protocol, defined in IEEE 802.1d, used to eliminate redundant data paths and to increase network efficiency. STP allows a network to have a topology that contains physical loops; it operates in bridges and switches. STP opens certain paths to create a tree topology, thereby preventing packets from looping endlessly on the network. To establish path redundancy, STP creates a tree that spans all of the switches in an extended network, forcing redundant paths into a standby, or blocked, state.

STP allows only one active path at a time between any two network devices (this prevents the loops) but establishes the redundant links as a backup if the initial link should fail. If STP costs change, or if one network segment in the STP becomes unreachable, the spanning tree algorithm reconfigures the STP topology and re-establishes the link by activating the standby path.

STPD

Spanning Tree Domain is an STP instance that contains one or more VLANs. The switch can run multiple STPDs, and each STPD has its own root bridge and active path. In the Extreme Networks implementation of STPD, each domain has a carrier VLAN (for carrying STP information) and one or more protected VLANs (for carrying the data).

syslog

A protocol used for the transmission of event notification messages across networks, originally developed on the University of California Berkeley Software Distribution (BSD) TCP/IP system implementations, and now embedded in many other operating systems and networked devices. A device generates a messages, a relay receives and forwards the messages, and a collector (a syslog server) receives the messages without relaying them.

syslog uses the UDP as its underlying transport layer mechanism. The UDP port that has been assigned to syslog is 514. (RFC 3164)

VLAN

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

VR-Control

This virtual router is part of the embedded system in Extreme Networks switches. VR-Control is used for internal communications between all the modules and subsystems in the switch. It has no ports, and you cannot assign any ports to it. It also cannot be associated with VLANs or routing protocols. (Referred to as VR-1 in earlier ExtremeXOS software versions.)

VR-Default

This virtual router is part of the embedded system in Extreme Networks switches. VR-Default is the default VR on the system. All data ports in the switch are assigned to this VR by default; you can add and delete ports from this VR. Likewise, VR-Default contains the default VLAN. Although you cannot delete the default VLAN from VR-Default, you can add and delete any user-created VLANs. One instance of each routing protocol is spawned for this VR, and they cannot be deleted. (Referred to as VR-2 in earlier ExtremeXOS software versions.)

VR-Mgmt

This virtual router is part of the embedded system in Extreme Networks switches. VR-Mgmt enables remote management stations to access the switch through Telnet, SSH, or SNMP sessions; and it owns the management port. The management port cannot be deleted from this VR, and no other ports can be added. The Mgmt VLAN is created VR-Mgmt, and it cannot be deleted; you cannot add or delete any other VLANs or any routing protocols to this VR. (Referred to as VR-0 in earlier ExtremeXOS software versions.)

VRRP

The Virtual Router Redundancy Protocol specifies an election protocol that dynamically assigns responsibility for a virtual router to one of the VRRP routers on a LAN. The VRRP router controlling the IP address(es) associated with a virtual router is called the master router, and forwards packets sent to these IP addresses. The election process provides dynamic failover in the forwarding responsibility should the master router become unavailable. In case the master router fails, the virtual IP address is mapped to a backup router's IP address; this backup becomes the master router. This allows any of the virtual router IP addresses on the LAN to be used as the default first-hop router by end-hosts. The advantage gained from using VRRP is a higher availability default path without requiring configuration of dynamic routing or router discovery protocols on every host. VRRP is defined in RFC 2338.

VXLAN

VXLAN is a network virtualization technology that leverages existing Layer 3 infrastructures to create tenant overlay networks. VXLAN addresses the requirements of a multi-tenant data center network infrastructure by:

- Increasing virtual network scalability to 16 million instances. This allows for tenant VLAN (Virtual LAN) isolation whereby multiple tenants can manage their own VLAN/VMAN and MAC address spaces.
- Adding an encapsulation that effectively hides VM MAC addresses from the physical network that results in physical networking devices to have smaller MAC and IP table sizes.

- Allowing for Layer 2 adjacency across IP networks by which DC network operators protect their investment in their current infrastructure. Additionally operators can distribute traffic loads across links efficiently using Layer 3 ECMP.

XNV

Extreme Network Virtualization is an ExtremeXOS feature that enables the software to support VM port movement, port configuration, and inventory on network switches.



Index

A

ACLs
 refreshing 2411
 smart refresh 2411
announcements 4, 5
ARP
 and VLAN aggregation 1507

C

CLI
 named components 11, 12

D

documentation
 feedback 4
Documentation, related 3, 4

E

EAPS
 names 11, 12

F

feedback 4
Flow Monitor 527–533, 1620–1622, 1705–1708, 1821, 1822,
2125–2127, 2664–2670

I

image
 primary and secondary 2034
inherit ports 1477, 2330
IPv6
 displaying 3286
ITU channels
 TDWDM XFP 1151

M

modular switch
 port number 13
MSTP
 identifiers 1401
 inherit ports 1477, 2330

N

names
 character types 11, 12, 1606, 1614, 1616, 1628, 1630, 1638,
1643–1647, 1659, 1669–1673, 1676, 1687
 conventions 11, 12, 1606, 1614, 1616, 1628, 1630, 1638,
1643–1647, 1659, 1669–1673, 1676, 1687
 maximum length of 11, 12, 1606, 1614, 1616, 1628, 1630, 1638,
1643–1647, 1659, 1669–1673, 1676, 1687
 VLAN, STP, EAPS 11, 12

O

Open Source Declaration 3, 4

P

partition 2034
port
 wildcard combinations 13
port-mirroring
 and ELSM 2208, 2212
 and load sharing 2208, 2209, 2212, 2213
 and sFlow 2208, 2209, 2212, 2213
 guidelines 845
 tagged and untagged frames 2209, 2210, 2212, 2213
 troubleshooting 2209, 2212
primary image 2034
product announcements 4, 5

R

refresh
 ACLs 2411

S

SCP2 2033
secondary image 2034
SFTP
 Using SFTP and SCP2 2033
smart refresh, ACLs 2411
software requirements for switches 8
STP
 inherit ports 1359, 1477, 2330
 names 11, 12
support, see technical support
switch
 software requirement 8
switch series, table 8

T

- TCP MD5 924
- technical support
 - contacting 4, 5
- troubleshooting
 - AppleTalk 1495, 1533
 - port mirroring 2210, 2213
- tunable DWDM XFP
 - ITU channels 1151

V

- virtual routers
 - commands 3554
- VLANs
 - names 11, 12
 - protocol-based 1495, 1533
- vMANs
 - names 11, 12

W

- wildcard combinations, port 13