



Junos[®] OS for EX Series Ethernet Switches

Interfaces Feature Guide for EX Series Switches

Release

15.1



Modified: 2016-05-22

Juniper Networks, Inc.
1133 Innovation Way
Sunnyvale, California 94089
USA
408-745-2000
www.juniper.net

Juniper Networks, Junos, Steel-Belted Radius, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. The Juniper Networks Logo, the Junos logo, and JunosE are trademarks of Juniper Networks, Inc. All other trademarks, service marks, registered trademarks, or registered service marks are the property of their respective owners.

Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

Junos[®] OS for EX Series Ethernet Switches Interfaces Feature Guide for EX Series Switches
Release 15.1
Copyright © 2016, Juniper Networks, Inc.
All rights reserved.

The information in this document is current as of the date on the title page.

YEAR 2000 NOTICE

Juniper Networks hardware and software products are Year 2000 compliant. Junos OS has no known time-related limitations through the year 2038. However, the NTP application is known to have some difficulty in the year 2036.

END USER LICENSE AGREEMENT

The Juniper Networks product that is the subject of this technical documentation consists of (or is intended for use with) Juniper Networks software. Use of such software is subject to the terms and conditions of the End User License Agreement ("EULA") posted at <http://www.juniper.net/support/eula.html>. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

Table of Contents

	About the Documentation	xv
	Documentation and Release Notes	xv
	Supported Platforms	xv
	Using the Examples in This Manual	xv
	Merging a Full Example	xvi
	Merging a Snippet	xvi
	Documentation Conventions	xvii
	Documentation Feedback	xix
	Requesting Technical Support	xix
	Self-Help Online Tools and Resources	xix
	Opening a Case with JTAC	xx
Part 1	Overview	
Chapter 1	Interfaces Overview	3
	EX Series Switches Interfaces Overview	3
	Network Interfaces	3
	Special Interfaces	4
	Understanding Interface Naming Conventions on EX Series Switches	6
	Physical Part of an Interface Name	6
	Logical Part of an Interface Name	8
	Wildcard Characters in Interface Names	8
	Understanding Aggregated Ethernet Interfaces and LACP	8
	Link Aggregation Group (LAG)	9
	Link Aggregation Control Protocol (LACP)	10
	Understanding Layer 3 Subinterfaces	11
	Understanding Unicast RPF	12
	Unicast RPF for Switches Overview	12
	Unicast RPF Implementation	13
	Unicast RPF Packet Filtering	13
	Bootstrap Protocol (BOOTP) and DHCP Requests	13
	Default Route Handling	13
	When to Enable Unicast RPF	13
	When Not to Enable Unicast RPF	14
	Limitations of the Unicast RPF Implementation on EX3200, EX4200, and EX4300 Switches	15
	Understanding IP Directed Broadcast	16
	IP Directed Broadcast Overview	16
	IP Directed Broadcast Implementation	16
	When to Enable IP Directed Broadcast	17
	When Not to Enable IP Directed Broadcast	17

	Understanding Interface Ranges on EX Series Switches	17
	802.1Q VLANs Overview	19
	Understanding Generic Routing Encapsulation	20
	Overview of GRE	20
	GRE Tunneling	21
	Encapsulation and De-Encapsulation on the Switch	21
	Number of Source and Destination Tunnels Allowed on a Switch	21
	Class of Service on GRE Tunnels	22
	Applying Firewall Filters to GRE Traffic	22
	Using a Firewall Filter to De-encapsulate GRE Traffic on a QFX5100 and OCX Series Switches	23
	Configuration Limitations	23
	Understanding Multicast Load Balancing on Aggregated 10-Gigabit Links for Routed Multicast Traffic on EX8200 Switches	24
	Create LAGs for Multicasting in Increments of 10 Gigabits	24
	When Should I Use Multicast Load Balancing?	25
	How Does Multicast Load Balancing Work?	26
	How Do I Implement Multicast Load Balancing on an EX8200 Switch?	27
	Understanding How Energy Efficient Ethernet Reduces Power Consumption on Interfaces	27
Part 2	Configuration	
Chapter 2	Configuration Examples	31
	Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch	31
	Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch	37
	Example: Configuring Layer 3 Subinterfaces for a Distribution Switch and an Access Switch	42
	Example: Configuring Unicast RPF on an EX Series Switch	49
	Example: Configuring IP Directed Broadcast on a Switch	54
	Example: Configuring Multicast Load Balancing for Use with Aggregated 10-Gigabit Ethernet Interfaces on EX8200 Switches	57
Chapter 3	Configuration Tasks	63
	Configuring Gigabit Ethernet Interfaces (CLI Procedure)	64
	Configuring VLAN Options and Port Mode	65
	Configuring the Link Settings	66
	Configuring the IP Options	68
	Configuring Gigabit Ethernet Interfaces (J-Web Procedure)	68
	Port Role Configuration with the J-Web Interface (with CLI References)	75
	Adding a Logical Unit Description to the Configuration	79
	Disabling a Physical Interface	80
	Disabling a Physical Interface	80
	Example: Disabling a Physical Interface	80
	Effect of Disabling Interfaces on T series PICs	81
	Disabling a Logical Interface	82

Configuring Flow Control	82
Configuring the Interface Address	84
Configuring the Interface Bandwidth	86
Configuring the Media MTU	87
Setting the Protocol MTU	88
Interface Ranges	89
Configuring Interface Ranges	89
Expanding Interface Range Member and Member Range Statements	93
Configuration Inheritance for Member Interfaces	94
Member Interfaces Inheriting Configuration from Configuration Groups	95
Interfaces Inheriting Common Configuration	96
Configuring Inheritance Range Priorities	97
Configuration Expansion Where Interface Range Is Used	97
Configuring Accounting for the Physical Interface	98
Accounting Profiles Overview	98
Configuring Accounting for the Physical Interface	99
Displaying Accounting Profile for the Physical Interface	100
Configuring Accounting for the Logical Interface	101
Accounting Profiles Overview	101
Configuring Accounting for the Logical Interface	101
Displaying Accounting Profile for the Logical Interface	102
Configuring Ethernet Loopback Capability	103
Configuring Gratuitous ARP	104
Configuring Static ARP Table Entries For Mapping IP Addresses to MAC Addresses	105
Disabling the Transmission of Redirect Messages on an Interface	106
Configuring Restricted and Unrestricted Proxy ARP	107
Enabling or Disabling SNMP Notifications on Logical Interfaces	108
Enabling or Disabling SNMP Notifications on Physical Interfaces	108
Configuring Aggregated Ethernet Links (CLI Procedure)	109
Configuring Aggregated Ethernet Interfaces (J-Web Procedure)	110
Configuring Aggregated Ethernet LACP (CLI Procedure)	113
Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure)	114
Configuring LACP Link Protection for a Single Link at the Global Level	116
Configuring LACP Link Protection for a Single Link at the Aggregated Interface Level	116
Configuring Subgroup Bundles to Provide LACP Link Protection to Multiple Links in an Aggregated Ethernet Interface	117
Configuring Aggregated Ethernet Link Protection	118
Configuring Link Protection for Aggregated Ethernet Interfaces	119
Configuring Primary and Backup Links for Link Aggregated Ethernet Interfaces	119
Reverting Traffic to a Primary Link When Traffic is Passing Through a Backup Link	119
Disabling Link Protection for Aggregated Ethernet Interfaces	119
Configuring Aggregated Ethernet Link Speed	120
Configuring Aggregated Ethernet Minimum Links	121
Configuring Tagged Aggregated Ethernet Interfaces	122

	Configuring a Layer 3 Subinterface (CLI Procedure)	123
	Configuring Unicast RPF (CLI Procedure)	123
	Disabling Unicast RPF (CLI Procedure)	125
	Configuring IP Directed Broadcast (CLI Procedure)	126
	Tracing Operations of an Individual Router or Switch Interface	127
	Tracing Operations of the Interface Process	128
	Setting the Mode on an SFP+ or SFP+ MACSec Uplink Module (CLI Procedure)	129
	Configuring the Media Type on Dual-Purpose Uplink Ports (CLI Procedure)	130
	Configuring Generic Routing Encapsulation Tunneling (CLI Procedure)	130
	Configuring a GRE Tunnel Port	131
	Configuring Tunnels to Use Generic Routing Encapsulation	131
	Configuring Multicast Load Balancing for Use with Aggregated 10-Gigabit Ethernet Links on EX8200 Switches (CLI Procedure)	132
	Configuring Energy Efficient Ethernet on Interfaces (CLI Procedure)	133
	Enabling EEE on an EEE-Capable Base-T Copper Ethernet Port	134
	Disabling EEE on a Base-T Copper Ethernet Port	134
	Damping Shorter Physical Interface Transitions	135
	Configuring Reflective Relay	136
Chapter 4	Configuration Statements	137
	[edit chassis] Configuration Statement Hierarchy on EX Series Switches	139
	Supported Statements in the [edit chassis] Hierarchy Level	140
	[edit interfaces] Configuration Statement Hierarchy on EX Series Switches	141
	[edit interfaces ae] Configuration Statement Hierarchy on EX Series Switches	142
	Supported Statements in the [edit interfaces ae] Hierarchy Level	142
	Unsupported Statements in the [edit interfaces ae] Hierarchy Level	145
	[edit interfaces ge] Configuration Statement Hierarchy on EX Series Switches	146
	Supported Statements in the [edit interfaces ge] Hierarchy Level	146
	Unsupported Statements in the [edit interfaces ge] Hierarchy Level	149
	[edit interfaces interface-range] Configuration Statement Hierarchy on EX Series Switches	149
	Supported Statements in the [edit interfaces interface-range] Hierarchy Level	150
	Unsupported Statements in the [edit interfaces interface-range] Hierarchy Level	154
	[edit interfaces lo] Configuration Statement Hierarchy on EX Series Switches	156
	Supported Statements in the [edit interfaces lo] Hierarchy Level	157
	Unsupported Statements in the [edit interfaces lo] Hierarchy Level	159
	[edit interfaces me] Configuration Statement Hierarchy on EX Series Switches	159
	Supported Statements in the [edit interfaces me] Hierarchy Level	160
	Unsupported Statements in the [edit interfaces me] Hierarchy Level	162

[edit interfaces vlan] Configuration Statement Hierarchy on EX Series	
Switches	163
Supported Statements in the [edit interfaces vlan] Hierarchy Level	163
Unsupported Statements in the [edit interfaces vlan] Hierarchy Level	166
[edit interfaces vme] Configuration Statement Hierarchy on EX Series	
Switches	166
Supported Statements in the [edit interfaces vme] Hierarchy Level	166
Unsupported Statements in the [edit interfaces vme] Hierarchy Level	169
[edit interfaces xe] Configuration Statement Hierarchy on EX Series	
Switches	169
Supported Statements in the [edit interfaces xe] Hierarchy Level	169
Unsupported Statements in the [edit interfaces xe] Hierarchy Level	173
[edit protocols lacp] Configuration Statement Hierarchy on EX Series	
Switches	173
Supported Statements in the [edit protocols lacp] Hierarchy Level	173
Unsupported Statements in the [edit protocols lacp] Hierarchy Level	174
802.3ad	174
accounting-profile	175
address	176
aggregated-devices	179
aggregated-ether-options	180
alarm (optics-options)	181
arp (Interfaces)	182
auto-negotiation	183
bandwidth (Interfaces)	185
broadcast	186
chassis	187
description	188
destination (Tunnels)	189
device-count	190
disable (Interface)	191
disable (Link Protection)	192
disable (Multicast Load Balancing)	192
eui-64	193
ether-options	194
ethernet (Aggregated Devices)	195
family (for EX Series switches)	196
filter	200
flow-control	201
force-up	202
gratuitous-arp-reply	202
hold-time (Physical Interface)	203
ieee-802-3az-eee	204
interface-range	205
interfaces (for EX Series switches)	207
lacp (Aggregated Ethernet)	215
lacp (802.3ad)	217
link-mode	218
link-protection	219

link-protection-sub-group (802.3ad)	220
link-protection-sub-group (aggregated-ether-options)	221
link-speed (Aggregated Ethernet)	222
loopback (Aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet)	224
mdi-mode	225
media-type (Dual-Purpose Uplink Ports)	226
member (Interface Ranges)	227
member-range	228
members	229
minimum-links	231
mtu	232
multicast-loadbalance	236
native-vlan-id	237
nd6-stale-time	238
no-redirects	239
non-revertive (Chassis)	239
non-revertive (Interfaces)	240
optics-options	241
periodic	242
pic	243
pic-mode	244
port-priority	245
port-mode	246
preferred	247
primary (Address on Interface)	248
proxy-arp	249
reflective-relay	250
rpf-check	251
sfpplus	252
source	253
speed	254
system-priority	255
system-priority	256
targeted-broadcast	257
traceoptions (Individual Interfaces)	258
traceoptions (Interface Process)	260
traps	262
ttl	263
tunnel	263
tunnel-port	264
unit	265
vlan (802.1Q Tagging)	266
vlan-id (VLAN Tagging and Layer 3 Subinterfaces)	267
vlan-tagging	268
warning	269
wavelength	270

Part 3	Administration	
Chapter 5	Routine Monitoring	277
	Monitoring Interface Status and Traffic	277
	Verifying the Status of a LAG Interface	279
	Verifying That LACP Is Configured Correctly and Bundle Members Are Exchanging	
	LACP Protocol Packets	279
	Verifying the LACP Setup	279
	Verifying That LACP Packets Are Being Exchanged	280
	Verifying That Layer 3 Subinterfaces Are Working	281
	Verifying Unicast RPF Status	281
	Verifying IP Directed Broadcast Status	284
	Verifying That Generic Routing Encapsulation Tunneling Is Working Correctly	284
	Verifying That EEE Is Saving Energy on Configured Ports	284
Chapter 6	Operational Commands	287
	Common Output Fields Description	287
	Damping Field	287
	Destination Class Field	288
	Enabled Field	288
	Filters Field	288
	Flags Fields	289
	Addresses, Flags Field	289
	Device Flags Field	289
	Family Flags Field	290
	Interface Flags Field	291
	Link Flags Field	292
	Logical Interface Flags Field	292
	Label-Switched Interface Traffic Statistics Field	292
	Policer Field	293
	Protocol Field	294
	RPF Failures Field	294
	Source Class Field	295
	monitor interface	296
	request diagnostics tdr	308
	show diagnostics tdr	310
	show interfaces (Aggregated Ethernet)	315
	show interfaces (GRE)	326
	show interfaces diagnostics optics	335
	show interfaces ge-	349
	show interfaces me0	361
	show interfaces xe-	368
	show interfaces queue	382
	show interfaces vlan	388
	show lacp interfaces	400
	show virtual-chassis vc-port diagnostics optics	405
	test interface restart-auto-negotiation	419

Part 4	Troubleshooting	
Chapter 7	Troubleshooting Procedures	423
	Troubleshooting an Aggregated Ethernet Interface	423
	Show Interfaces Command Shows the LAG is Down	423
	Logical Interface Statistics Do Not Reflect All Traffic	424
	IPv6 Interface Traffic Statistics Are Not Supported	424
	SNMP Counters ifHCInBroadcastPkts and ifInBroadcastPkts Are Always 0	424
	Troubleshooting Network Interfaces on EX3200 Switches	424
	The interface on one of the last four built-in network ports in an EX3200 switch (for example, interface ge-0/0/23) is down	425
	The interface on the port in which an SFP or SFP+ transceiver is installed in an SFP+ uplink module is down	425
	Troubleshooting Network Interfaces on EX4200 Switches	426
	The interface on the port in which an SFP or SFP+ transceiver is installed is down	426
	Troubleshooting Uplink Module Installation or Replacement on EX3200 Switches	427
	One of the last four network ports on an EX3200 switch with an SFP or SFP+ uplink module installed is disabled	427
	A port on an SFP uplink module installed in an EX3200 switch is disabled	427
	Troubleshooting Interface Configuration and Cable Faults	428
	Interface Configuration or Connectivity Is Not Working	428
	Troubleshooting Unicast RPF	429
	Legitimate Packets Are Discarded	429
	Diagnosing a Faulty Twisted-Pair Cable (CLI Procedure)	430

List of Figures

Part 1	Overview	
Chapter 1	Interfaces Overview	3
	Figure 1: Symmetrically Routed Interfaces	14
	Figure 2: Asymmetrically Routed Interfaces	15
	Figure 3: 40-Gigabit LAGs on EX8200 Switches	25
Part 2	Configuration	
Chapter 2	Configuration Examples	31
	Figure 4: Topology for LAGs Connecting an EX4200 Virtual Chassis Access Switch to an EX4200 Virtual Chassis Distribution Switch	33
	Figure 5: Topology for IP Directed Broadcast	55
	Figure 6: 40-Gigabit LAG Composed of Four 10-Gigabit Links	59

List of Tables

	About the Documentation	xv
	Table 1: Notice Icons	xvii
	Table 2: Text and Syntax Conventions	xvii
Part 1	Overview	
Chapter 1	Interfaces Overview	3
	Table 3: Network Interface Types and Purposes	4
	Table 4: Special Interface Types and Purposes	4
	Table 5: Maximum Interfaces per LAG and Maximum LAGs per Switch	9
	Table 6: Firewall Filter Application Points for Tunneled Packets	22
	Table 7: Features Not Supported with GRE	23
	Table 8: Hashing Algorithms Used by Multicast Load Balancing	26
Part 2	Configuration	
Chapter 2	Configuration Examples	31
	Table 9: Components of the Topology for Connecting a Virtual Chassis Access Switch to a Virtual Chassis Distribution Switch	33
	Table 10: Components of the Topology for Creating Layer 3 Subinterfaces on an Access Switch and a Distribution Switch	43
	Table 11: Components of the IP Directed Broadcast Topology	55
Chapter 3	Configuration Tasks	63
	Table 12: Port Edit Options	71
	Table 13: Recommended CoS Settings for Port Roles	74
	Table 14: Port Role Configuration Summary	75
	Table 15: Recommended CoS Settings for Port Roles	78
	Table 16: Effect of set interfaces disable <interface_name> on T series PICs	81
	Table 17: Aggregated Ethernet Interface Options	111
	Table 18: VLAN Options	112
	Table 19: IP Options	112
Chapter 4	Configuration Statements	137
	Table 20: Unsupported [edit interfaces ae] Configuration Statements on EX Series Switches	145
	Table 21: Unsupported [edit interfaces interface-range] Configuration Statements for EX Series Switches	154
	Table 22: Unsupported [edit interfaces lo] Configuration Statements for EX Series Switches	159
	Table 23: Unsupported [edit interfaces me] Configuration Statements for EX Series Switches	162

Table 24: Protocol Families and Supported Interface Types	198
Table 25: Interface Types and Their Supported Protocol Families	213

Part 3

Chapter 6

Administration

Operational Commands	287
---------------------------------------	------------

Table 26: Output Control Keys for the monitor interface interface-name Command	296
Table 27: Output Control Keys for the monitor interface traffic Command	297
Table 28: monitor interface Output Fields	298
Table 29: request diagnostics tdr Output Fields	309
Table 30: show diagnostics tdr Output Fields	311
Table 31: Aggregated Ethernet show interfaces Output Fields	315
Table 32: GRE show interfaces Output Fields	327
Table 33: show interfaces diagnostics optics Output Fields	335
Table 34: show interfaces ge- Output Fields	350
Table 35: show interfaces me0 Output Fields	361
Table 36: show interfaces xe- Output Fields	369
Table 37: show interfaces queue Output Fields	382
Table 38: show interfaces vlan Output Fields	389
Table 39: show lacp interfaces Output Fields	401
Table 40: show virtual-chassis vc-port diagnostics optics Output Fields	406

About the Documentation

- Documentation and Release Notes on page xv
- Supported Platforms on page xv
- Using the Examples in This Manual on page xv
- Documentation Conventions on page xvii
- Documentation Feedback on page xix
- Requesting Technical Support on page xix

Documentation and Release Notes

To obtain the most current version of all Juniper Networks® technical documentation, see the product documentation page on the Juniper Networks website at <http://www.juniper.net/techpubs/>.

If the information in the latest release notes differs from the information in the documentation, follow the product Release Notes.

Juniper Networks Books publishes books by Juniper Networks engineers and subject matter experts. These books go beyond the technical documentation to explore the nuances of network architecture, deployment, and administration. The current list can be viewed at <http://www.juniper.net/books>.

Supported Platforms

For the features described in this document, the following platforms are supported:

- EX Series

Using the Examples in This Manual

If you want to use the examples in this manual, you can use the **load merge** or the **load merge relative** command. These commands cause the software to merge the incoming configuration into the current candidate configuration. The example does not become active until you commit the candidate configuration.

If the example configuration contains the top level of the hierarchy (or multiple hierarchies), the example is a *full example*. In this case, use the **load merge** command.

If the example configuration does not start at the top level of the hierarchy, the example is a *snippet*. In this case, use the **load merge relative** command. These procedures are described in the following sections.

Merging a Full Example

To merge a full example, follow these steps:

1. From the HTML or PDF version of the manual, copy a configuration example into a text file, save the file with a name, and copy the file to a directory on your routing platform.

For example, copy the following configuration to a file and name the file **ex-script.conf**. Copy the **ex-script.conf** file to the **/var/tmp** directory on your routing platform.

```
system {
  scripts {
    commit {
      file ex-script.xml;
    }
  }
}
interfaces {
  fxp0 {
    disable;
    unit 0 {
      family inet {
        address 10.0.0.1/24;
      }
    }
  }
}
```

2. Merge the contents of the file into your routing platform configuration by issuing the **load merge** configuration mode command:

```
[edit]
user@host# load merge /var/tmp/ex-script.conf
load complete
```

Merging a Snippet

To merge a snippet, follow these steps:

1. From the HTML or PDF version of the manual, copy a configuration snippet into a text file, save the file with a name, and copy the file to a directory on your routing platform.

For example, copy the following snippet to a file and name the file **ex-script-snippet.conf**. Copy the **ex-script-snippet.conf** file to the **/var/tmp** directory on your routing platform.

```
commit {
  file ex-script-snippet.xml; }
```

2. Move to the hierarchy level that is relevant for this snippet by issuing the following configuration mode command:

```
[edit]
user@host# edit system scripts
[edit system scripts]
```

3. Merge the contents of the file into your routing platform configuration by issuing the **load merge relative** configuration mode command:

```
[edit system scripts]
user@host# load merge relative /var/tmp/ex-script-snippet.conf
load complete
```

For more information about the **load** command, see the *CLI User Guide*.

Documentation Conventions

Table 1 on page xvii defines notice icons used in this guide.

Table 1: Notice Icons







Icon	Meaning	Description
	Informational note	Indicates important features or instructions.
	Caution	Indicates a situation that might result in loss of data or hardware damage.
	Warning	Alerts you to the risk of personal injury or death.
	Laser warning	Alerts you to the risk of personal injury from a laser.
	Tip	Indicates helpful information.
	Best practice	Alerts you to a recommended use or implementation.

Table 2 on page xvii defines the text and syntax conventions used in this guide.

Table 2: Text and Syntax Conventions

Convention	Description	Examples
Bold text like this	Represents text that you type.	To enter configuration mode, type the configure command: user@host> configure

Table 2: Text and Syntax Conventions (*continued*)

Convention	Description	Examples
Fixed-width text like this	Represents output that appears on the terminal screen.	<code>user@host> show chassis alarms</code> <code>No alarms currently active</code>
<i>Italic text like this</i>	<ul style="list-style-type: none">Introduces or emphasizes important new terms.Identifies guide names.Identifies RFC and Internet draft titles.	<ul style="list-style-type: none">A policy <i>term</i> is a named structure that defines match conditions and actions.<i>Junos OS CLI User Guide</i>RFC 1997, <i>BGP Communities Attribute</i>
<i>Italic text like this</i>	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name: [edit] root@# set system domain-name <i>domain-name</i>
Text like this	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	<ul style="list-style-type: none">To configure a stub area, include the stub statement at the [edit protocols ospf area area-id] hierarchy level.The console port is labeled CONSOLE.
< > (angle brackets)	Encloses optional keywords or variables.	stub <default-metric <i>metric</i>>;
(pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	broadcast multicast (<i>string1</i> <i>string2</i> <i>string3</i>)
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	rsvp { # Required for dynamic MPLS only
[] (square brackets)	Encloses a variable for which you can substitute one or more values.	community name members [<i>community-ids</i>]
Indentation and braces ({ })	Identifies a level in the configuration hierarchy.	[edit] routing-options { static { route default { nexthop <i>address</i> ; retain; } } }
;(semicolon)	Identifies a leaf statement at a configuration hierarchy level.	
GUI Conventions		
Bold text like this	Represents graphical user interface (GUI) items you click or select.	<ul style="list-style-type: none">In the Logical Interfaces box, select All Interfaces.To cancel the configuration, click Cancel.

Table 2: Text and Syntax Conventions (*continued*)

Convention	Description	Examples
> (bold right angle bracket)	Separates levels in a hierarchy of menu selections.	In the configuration editor hierarchy, select Protocols>Ospf .

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can provide feedback by using either of the following methods:

- Online feedback rating system—On any page of the Juniper Networks TechLibrary site at <http://www.juniper.net/techpubs/index.html>, simply click the stars to rate the content, and use the pop-up form to provide us with information about your experience. Alternately, you can use the online feedback form at <http://www.juniper.net/techpubs/feedback/>.
- E-mail—Send your comments to techpubs-comments@juniper.net. Include the document or topic name, URL or page number, and software version (if applicable).

Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active J-Care or Partner Support Service support contract, or are covered under warranty, and need post-sales technical support, you can access our tools and resources online or open a case with JTAC.

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the *JTAC User Guide* located at <http://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf>.
- Product warranties—For product warranty information, visit <http://www.juniper.net/support/warranty/>.
- JTAC hours of operation—The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

Self-Help Online Tools and Resources

For quick and easy problem resolution, Juniper Networks has designed an online self-service portal called the Customer Support Center (CSC) that provides you with the following features:

- Find CSC offerings: <http://www.juniper.net/customers/support/>
- Search for known bugs: <http://www2.juniper.net/kb/>
- Find product documentation: <http://www.juniper.net/techpubs/>
- Find solutions and answer questions using our Knowledge Base: <http://kb.juniper.net/>

- Download the latest versions of software and review release notes:
<http://www.juniper.net/customers/csc/software/>
- Search technical bulletins for relevant hardware and software notifications:
<http://kb.juniper.net/InfoCenter/>
- Join and participate in the Juniper Networks Community Forum:
<http://www.juniper.net/company/communities/>
- Open a case online in the CSC Case Management tool: <http://www.juniper.net/cm/>

To verify service entitlement by product serial number, use our Serial Number Entitlement (SNE) Tool: <https://tools.juniper.net/SerialNumberEntitlementSearch/>

Opening a Case with JTAC

You can open a case with JTAC on the Web or by telephone.

- Use the Case Management tool in the CSC at <http://www.juniper.net/cm/>.
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

For international or direct-dial options in countries without toll-free numbers, see <http://www.juniper.net/support/requesting-support.html>.

PART 1

Overview

- [Interfaces Overview on page 3](#)

CHAPTER 1

Interfaces Overview

- [EX Series Switches Interfaces Overview on page 3](#)
- [Understanding Interface Naming Conventions on EX Series Switches on page 6](#)
- [Understanding Aggregated Ethernet Interfaces and LACP on page 8](#)
- [Understanding Layer 3 Subinterfaces on page 11](#)
- [Understanding Unicast RPF on page 12](#)
- [Understanding IP Directed Broadcast on page 16](#)
- [Understanding Interface Ranges on EX Series Switches on page 17](#)
- [802.1Q VLANs Overview on page 19](#)
- [Understanding Generic Routing Encapsulation on page 20](#)
- [Understanding Multicast Load Balancing on Aggregated 10-Gigabit Links for Routed Multicast Traffic on EX8200 Switches on page 24](#)
- [Understanding How Energy Efficient Ethernet Reduces Power Consumption on Interfaces on page 27](#)

EX Series Switches Interfaces Overview

Juniper Networks EX Series Ethernet Switches have two types of interfaces: network interfaces and special interfaces. This topic provides brief information about these interfaces. For additional information, see the *[Junos OS Interfaces Fundamentals Configuration Guide](#)*.

For information about interface-naming conventions on EX Series switches, see “[Understanding Interface Naming Conventions on EX Series Switches](#)” on page 6.

This topic describes:

- [Network Interfaces on page 3](#)
- [Special Interfaces on page 4](#)

Network Interfaces

Network interfaces connect to the network and carry network traffic. [Table 3 on page 4](#) lists the types of network interfaces supported on EX Series switches.

Table 3: Network Interface Types and Purposes

Type	Purpose
Aggregated Ethernet interfaces	All EX Series switches allow you to group Ethernet interfaces at the physical layer to form a single link layer interface, also known as a <i>link aggregation group (LAG)</i> or <i>bundle</i> . These aggregated Ethernet interfaces help to balance traffic and increase the uplink bandwidth.
LAN access interfaces	Use these EX Series switch interfaces to connect a personal computer, laptop, file server, or printer to the network. When you power on an EX Series switch and use the factory-default configuration, the software automatically configures interfaces in access mode for each of the network ports. The default configuration also enables autonegotiation for both speed and link mode.
Power over Ethernet (PoE) interfaces	EX Series switches provide PoE network ports with various switch models. These ports can be used to connect voice over IP (VoIP) telephones, wireless access points, video cameras, and point-of-sale devices to safely receive power from the same access ports that are used to connect personal computers to the network. PoE interfaces are enabled by default in the factory configuration.
Trunk interfaces	EX Series access switches can be connected to a distribution switch or customer-edge (CE) switches or routers. To use a port for this type of connection, you must explicitly configure the network interface for trunk mode. The interfaces from the distribution switch or CE switch to the access switches must also be configured for trunk mode.

Special Interfaces

Table 4 on page 4 lists the types of special interfaces supported on EX Series switches.

Table 4: Special Interface Types and Purposes

Type	Purpose
Console port	Each EX Series switch has a serial port, labeled CON or CONSOLE , for connecting tty-type terminals to the switch using standard PC-type tty cables. The console port does not have a physical address or IP address associated with it. However, it is an interface in the sense that it provides access to the switch. On an EX3300 Virtual Chassis, an EX4200 Virtual Chassis, or an EX4500 Virtual Chassis, you can access the master and configure all members of the Virtual Chassis through any member's console port. For more information about the console port in a Virtual Chassis, see <i>Understanding Global Management of a Virtual Chassis</i> .
Loopback	All EX Series switches have this software-only virtual interface that is always up. The loopback interface provides a stable and consistent interface and IP address on the switch.
Management interface	The Juniper Networks Junos operating system (Junos OS) for EX Series switches automatically creates the switch's management Ethernet interface, me0 . The management Ethernet interface provides an out-of-band method for connecting to the switch. To use me0 as a management port, you must configure its logical port, me0.0 , with a valid IP address. You can connect to the management interface over the network using utilities such as SSH or Telnet. SNMP can use the management interface to gather statistics from the switch. (The management interface me0 is analogous to the fxp0 interfaces on routers running Junos OS.)

Table 4: Special Interface Types and Purposes (*continued*)

Type	Purpose
Integrated Routing and Bridging (IRB) Interface or Routed VLAN Interface (RVI)	<p>EX Series switches use an integrated routing and bridging (IRB) interface or Routed VLAN Interface (RVI) to route traffic from one broadcast domain to another and to perform other Layer 3 functions such as traffic engineering. These functions are typically performed by a router interface in a traditional network.</p> <p>The IRB interface or RVI functions as a logical router, eliminating the need for having both a switch and a router. These interfaces must be configured as part of a broadcast domain or virtual private LAN service (VPLS) routing instance for Layer 3 traffic to be routed from.</p>
Virtual Chassis port (VCP) interfaces	<p>Virtual Chassis ports (VCPs) are used to interconnect switches in a Virtual Chassis:</p> <ul style="list-style-type: none"> EX3300 switches—Port 2 and port 3 of the SFP+ uplink ports are preconfigured as VCPs and can be used to interconnect up to six EX3300 switches in an EX3300 Virtual Chassis. See <i>Setting an Uplink Port on an EX Series Switch as a Virtual Chassis Port (CLI Procedure)</i>. EX4200 and EX4500 switches—Each EX4200 switch or each EX4500 switch with a Virtual Chassis module installed has two dedicated VCPs on its rear panel. These ports can be used to interconnect up to ten EX4200 switches in an EX4200 Virtual Chassis, up to ten EX4500 switches in an EX4500 Virtual Chassis, and up to ten switches in a mixed EX4200 and EX4500 Virtual Chassis. When you power on switches that are interconnected in this manner, the software automatically configures the VCP interfaces for the dedicated ports that have been interconnected. These VCP interfaces are not configurable or modifiable. See <i>Understanding the High-Speed Interconnection of the Dedicated Virtual Chassis Ports Connecting EX4200, EX4500, and EX4550 Member Switches</i>. <p>You can also interconnect EX4200 and EX4500 switches by using uplink module ports. Using uplink ports allows you to connect switches over longer distances than you can by using the dedicated VCPs. To use the uplink ports as VCPs, you must explicitly configure the uplink module ports on the members you want to connect as VCPs. See <i>Setting an Uplink Port on an EX Series Switch as a Virtual Chassis Port (CLI Procedure)</i> or <i>Setting an Uplink Port as a Virtual Chassis Port on an EX4500 or EX4550 Switch (CLI Procedure)</i>.</p> <ul style="list-style-type: none"> EX4300 switches—All QSFP+ ports are configured as VCPs, by default. See <i>Understanding EX4300 Virtual Chassis</i>. <p>You can also interconnect EX4300 switches into a Virtual Chassis by using SFP+ uplink module ports as VCPs. Using uplink ports as VCPs allows you to connect switches over longer distances than you can by using the QSFP+ ports as VCPs. To use the uplink ports as VCPs, you must explicitly configure the uplink module ports on the members you want to connect as VCPs. See <i>Setting an Uplink Port on an EX Series Switch as a Virtual Chassis Port (CLI Procedure)</i>.</p> <ul style="list-style-type: none"> EX8200 switches—EX8200 switches can be connected to an XRE200 External Routing Engine to create an EX8200 Virtual Chassis. The XRE200 External Routing Engine has dedicated VCPs that connect to ports on the internal Routing Engines of the EX8200 switches and can connect to another XRE200 External Routing Engine for redundancy. These ports require no configuration. <p>You can also connect two members of an EX8200 Virtual Chassis so that they can exchange Virtual Chassis Control Protocol (VCCP) traffic. To do so, you explicitly configure network ports on the EX8200 switches as VCPs. See <i>Understanding Virtual Chassis Ports in an EX8200 Virtual Chassis</i>.</p>
Virtual management Ethernet (VME) interface	<p>EX3300, EX4200, EX4300, and EX4500 switches have a VME interface. This is a logical interface that is used for Virtual Chassis configurations and allows you to manage all the members of the Virtual Chassis through the master. For more information about the VME interface, see <i>Understanding Global Management of a Virtual Chassis</i>.</p> <p>EX8200 switches do not use a VME interface. An EX8200 Virtual Chassis is managed through the management Ethernet (me0) interface on the XRE200 External Routing Engine.</p>

Related Documentation

- [EX2200 Switches Hardware Overview](#)
- [EX3200 Switches Hardware Overview](#)
- [EX3300 Switches Hardware Overview](#)
- [EX4200 Switches Hardware Overview](#)
- [EX4300 Switches Hardware Overview](#)
- [EX4500 Switches Hardware Overview](#)
- [EX6210 Switch Hardware Overview](#)
- [EX8208 Switch Hardware Overview](#)
- [EX8216 Switch Hardware Overview](#)
- [XRE200 External Routing Engine Hardware Overview](#)
- [Understanding PoE on EX Series Switches](#)
- [Understanding Aggregated Ethernet Interfaces and LACP on page 8](#)
- [Understanding Layer 3 Subinterfaces on page 11](#)

Understanding Interface Naming Conventions on EX Series Switches

Juniper Networks EX Series Ethernet Switches use a naming convention for defining the interfaces that is similar to that of other platforms running under Juniper Networks Junos operating system (Junos OS). This topic provides brief information about the naming conventions used for interfaces on EX Series switches. For additional information, see the [Junos OS Network Interfaces Configuration Guide](#).

This topic describes:

- [Physical Part of an Interface Name on page 6](#)
- [Logical Part of an Interface Name on page 8](#)
- [Wildcard Characters in Interface Names on page 8](#)

Physical Part of an Interface Name

Network interfaces in Junos OS are specified as follows:

type-fpc / pic / port

EX Series switches apply this convention as follows:

- *type*—EX Series interfaces use the following media types:
 - **ge**—Gigabit Ethernet interface
 - **xe**—10 Gigabit Ethernet interface

- **et**—40 Gigabit Ethernet interface
- **fpc**—Flexible PIC Concentrator. EX Series interfaces use the following convention for the FPC number in interface names:
 - On an EX2200 switch, an EX3200 switch, a standalone EX3300 switch, a standalone EX4200 switch, a standalone EX4300 switch, a standalone EX4500, and a standalone EX4550 switch, FPC refers to the switch itself. The FPC number is **0** by default on these switches.
 - On an EX3300 Virtual Chassis, an EX4200 Virtual Chassis, an EX4300 Virtual Chassis, an EX4500 Virtual Chassis, an EX4550 Virtual Chassis, or a mixed Virtual Chassis, the FPC number indicates the member ID of the switch in the Virtual Chassis.
 - On an EX6200 switch and a standalone EX8200 switch, the FPC number indicates the slot number of the line card that contains the physical interface. On an EX6200 switch, the FPC number also indicates the slot number of the Switch Fabric and Routing Engine (SRE) module that contains the uplink port.
 - On an EX8200 Virtual Chassis, the FPC number indicates the slot number of the line card on the Virtual Chassis. The line card slots on Virtual Chassis member 0 are numbered 0 through 15; on member 1, they are numbered 16 through 31, and so on.
- **pic**—EX Series interfaces use the following convention for the PIC (Physical Interface Card) number in interface names:
 - On EX2200, EX3200, EX3300, EX4200, EX4500 switch, and EX4550 switches, the PIC number is **0** for all built-in interfaces (interfaces that are not uplink ports).
 - On EX2200, EX3200, EX3300, and EX4200 switches, the PIC number is **1** for uplink ports.
 - On EX4300 switches, the PIC number is **0** for built-in network ports, **1** for built-in QSFP+ ports (located on the rear panel of the switch), and **2** for uplink module ports.
 - On EX4500 switches, the PIC number is **1** for ports on the left-hand uplink module and **2** for ports on the right-hand uplink module.
 - On EX4550 switches, the PIC number is **1** for ports in the expansion module or Virtual Chassis module installed in the module slot on the front panel of the switch and **2** for those in the expansion module or Virtual Chassis module installed in the module slot on the rear panel of the switch.
 - On EX6200 and EX8200 switches, the PIC number is always **0**.
- **port**—EX Series interfaces use the following convention for port numbers:
 - On EX2200, EX3200, EX3300, EX4200, EX4300, EX4500, and EX4550 switches, built-in network ports are numbered from left to right. On models that have two rows of ports, the ports on the top row start with **0** followed by the remaining even-numbered ports, and the ports on the bottom row start with **1** followed by the remaining odd-numbered ports.
 - Uplink ports in EX2200, EX3200, EX3300, EX4200, EX4300, EX4500, and EX4550 switches are labeled from left to right, starting with **0**.

- On EX6200 and EX8200 switches, the network ports are numbered from left to right on each line card. On line cards that have two rows of ports, the ports on the top row start with 0 followed by the remaining even-numbered ports, and the ports on the bottom row start with 1 followed by the remaining odd-numbered ports.
- Uplink ports on an SRE module in an EX6200 switch are labeled from left to right, starting with 0.

Logical Part of an Interface Name

The logical unit part of the interface name corresponds to the logical unit number, which can be a number from 0 through 16384. In the virtual part of the name, a period (.) separates the port and logical unit numbers: *type-fpc/pic/port.logical-unit-number*. For example, if you issue the **show ethernet-switching interfaces** command on a system with a default VLAN, the resulting display shows the logical interfaces associated with the VLAN:

Interface	State	VLAN members	Blocking
ge-0/0/0.0	down	remote-analyzer	unblocked
ge-0/0/1.0	down	default	unblocked
ge-0/0/10.0	down	default	unblocked

Wildcard Characters in Interface Names

In the **show interfaces** and **clear interfaces** commands, you can use wildcard characters in the *interface-name* option to specify groups of interface names without having to type each name individually. You must enclose all wildcard characters except the asterisk (*) in quotation marks (" ").

Related Documentation

- [EX Series Switches Interfaces Overview on page 3](#)
- [Configuring Gigabit Ethernet Interfaces \(CLI Procedure\) on page 64](#)
- [Configuring Gigabit Ethernet Interfaces \(CLI Procedure\)](#)

Understanding Aggregated Ethernet Interfaces and LACP

IEEE 802.3ad link aggregation enables you to group Ethernet interfaces to form a single link layer interface, also known as a *link aggregation group (LAG)* or *bundle*.

Aggregating multiple links between physical interfaces creates a single logical point-to-point trunk link or a LAG. The LAG balances traffic across the member links within an aggregated Ethernet bundle and effectively increases the uplink bandwidth. Another advantage of link aggregation is increased availability, because the LAG is composed of multiple member links. If one member link fails, the LAG continues to carry traffic over the remaining links.

Link Aggregation Control Protocol (LACP), a component of IEEE 802.3ad, provides additional functionality for LAGs.

This topic describes:

- [Link Aggregation Group \(LAG\) on page 9](#)
- [Link Aggregation Control Protocol \(LACP\) on page 10](#)

Link Aggregation Group (LAG)

You configure a LAG by specifying the link number as a physical device and then associating a set of interfaces (ports) with the link. All the interfaces must have the same speed and be in full-duplex mode. Juniper Networks Junos operating system (Junos OS) for EX Series Ethernet Switches assigns a unique ID and port priority to each interface. The ID and priority are not configurable.

The number of interfaces that can be grouped into a LAG and the total number of LAGs supported on a switch varies according to switch model. [Table 5 on page 9](#) lists the EX Series switches and the maximum number of interfaces per LAG and the maximum number of LAGs they support. MX Series devices can support up to 64 LAGs.

Table 5: Maximum Interfaces per LAG and Maximum LAGs per Switch

Switch	Maximum Interfaces per LAG	Maximum LAGs
EX2200	8	32
EX3200	8	32
EX3300 and EX3300 Virtual Chassis	8	111
EX4200 and EX4200 Virtual Chassis	8	111
EX4300 and EX4300 Virtual Chassis	16	112
EX4500, EX4500 Virtual Chassis, EX4550, and EX4550 Virtual Chassis	8	111
EX6200	8	111
EX8200	12	255
EX8200 Virtual Chassis	12	239

When configuring LAGs, consider the following guidelines:

- You must configure the LAG on both sides of the link.
- You must set the interfaces on either side of the link to the same speed.
- You can configure and apply firewall filters on a LAG.

- You can optionally configure LACP for link negotiation.
- You can optionally configure LACP for link protection.

You can combine physical Ethernet ports belonging to different member switches of a Virtual Chassis configuration to form a LAG. See *Understanding EX Series Virtual Chassis Port Link Aggregation* and *Understanding Link Aggregation in an EX8200 Virtual Chassis*.



NOTE: The interfaces that are included within a LAG are sometimes referred to as *member interfaces*. Do not confuse this term with *member switches*, which refers to switches that are interconnected as a Virtual Chassis. It is possible to create a LAG that is composed of member interfaces that are located in different member switches of a Virtual Chassis.

A LAG hashing algorithm determines how traffic entering a LAG is placed onto the bundle's member links. The LAG hashing algorithm tries to manage bandwidth by evenly load-balancing all incoming traffic across the member links in the bundle. You can configure the fields used by the LAG hashing algorithm on some EX Series switches. See *Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic (CLI Procedure)*.

A LAG creates a single logical point-to-point connection. A typical deployment for a LAG would be to aggregate trunk links between an access switch and a distribution switch or customer edge (CE) router.

Link Aggregation Control Protocol (LACP)

When LACP is configured, it detects misconfigurations on the local end or the remote end of the link. Thus, LACP can help prevent communication failure:

- When LACP is not enabled, a local LAG might attempt to transmit packets to a remote single interface, which causes the communication to fail.
- When LACP is enabled, a local LAG cannot transmit packets unless a LAG with LACP is also configured on the remote end of the link.

By default, Ethernet links do not exchange LACP protocol data units (PDUs), which contain information about the state of the link. You can configure Ethernet links to actively transmit LACP PDUs, or you can configure the links to passively transmit them, sending out LACP PDUs only when the Ethernet link receives them from the remote end. The transmitting link is known as the *actor* and the receiving link is known as the *partner*.

In a scenario where a dual-homed server is deployed with a switch, the network interface cards form a LAG with the switch. During a server upgrade, the server might not be able to exchange LACP PDUs. In such a situation, you can configure an interface to be in the **up** state even if no PDUs are exchanged. Use the **force-up** statement to configure an interface when the peer has limited LACP capability. The interface selects the associated LAG by default, whether the switch and peer are both in active or passive mode. When PDUs are not received, the partner is considered to be working in the passive mode. Therefore, LACP PDU transmissions are controlled by the transmitting link.

If the remote end of the LAG link is a security device, LACP might not be supported because security devices require a deterministic configuration. In such a scenario, do not configure LACP. All links in the LAG are permanently operational unless the switch detects a link failure within the Ethernet physical layer or data link layers.

**Related
Documentation**

- [Configuring Aggregated Ethernet Links \(CLI Procedure\) on page 109](#)
- [Configuring Aggregated Ethernet LACP \(CLI Procedure\) on page 113](#)
- [Configuring LACP Link Protection of Aggregated Ethernet Interfaces \(CLI Procedure\) on page 114](#)
- [Junos OS Network Interfaces Configuration Guide](#)

Understanding Layer 3 Subinterfaces

A Layer 3 subinterface is a logical division of a physical interface that operates at the network level and therefore can receive and forward 802.1Q VLAN tags. You can use Layer 3 subinterfaces to route traffic among multiple VLANs along a single trunk line that connects a Juniper Networks EX Series Ethernet Switch to a Layer 2 switch. Only one physical connection is required between the switches. This topology is often called a *router on a stick* or a *one-armed router* when the Layer 3 device is a router.

To create Layer 3 subinterfaces on an EX Series switch, you enable VLAN tagging, partition the physical interface into logical partitions, and bind the VLAN ID to the logical interface.

You can partition one physical interface into up to 4094 different subinterfaces, one for each VLAN. We recommend that you use the VLAN ID as the subinterface number when you configure the subinterface. Juniper Networks Junos operating system (Junos OS) reserves VLAN IDs 0 and 4095.

VLAN tagging places the VLAN ID in the frame header, allowing each physical interface to handle multiple VLANs. When you configure multiple VLANs on an interface, you must also enable tagging on that interface. Junos OS on EX Series switches supports a subset of the 802.1Q standard for receiving and forwarding routed or bridged Ethernet frames with single VLAN tags and running Virtual Router Redundancy Protocol (VRRP) over 802.1Q-tagged interfaces. Double-tagging is not supported.

**Related
Documentation**

- [EX Series Switches Interfaces Overview on page 3](#)
- [Example: Configuring Layer 3 Subinterfaces for a Distribution Switch and an Access Switch on page 42](#)
- [Junos OS Ethernet Interfaces Configuration Guide](#)

Understanding Unicast RPF

Unicast reverse-path forwarding (RPF) helps protect the switch against denial-of-service (DoS) and distributed denial-of-service (DDoS) attacks by verifying the unicast source address of each packet that arrives on an ingress interface where unicast RPF is enabled. It also helps ensure that traffic arriving on ingress interfaces comes from a network source that the receiving interface can reach.

When you enable unicast RPF, by default the switch forwards a packet only if the receiving interface is the best return path to the packet's unicast source address. This is known as strict mode unicast RPF. You can also enable loose mode, which means that the system checks to see if the packet has a source address with a corresponding prefix in the routing table but does not check whether the receiving interface is the best return path to the packet's unicast source address.



.....

NOTE: On Juniper Networks EX3200, EX4200, and EX4300 Ethernet Switches, the switch applies unicast RPF *globally* to all interfaces when unicast RPF is configured on any interface. For additional information, see [“Limitations of the Unicast RPF Implementation on EX3200, EX4200, and EX4300 Switches” on page 15.](#)

.....

This topic covers:

- [Unicast RPF for Switches Overview on page 12](#)
- [Unicast RPF Implementation on page 13](#)
- [When to Enable Unicast RPF on page 13](#)
- [When Not to Enable Unicast RPF on page 14](#)
- [Limitations of the Unicast RPF Implementation on EX3200, EX4200, and EX4300 Switches on page 15](#)

Unicast RPF for Switches Overview

Unicast RPF functions as an ingress filter that reduces the forwarding of IP packets that might be spoofing an address. By default, unicast RPF is disabled on the switch interfaces.

The type of unicast RPF provided on the switches—that is, strict mode unicast RPF is especially useful on untrusted interfaces. An untrusted interface is an interface where untrusted users or processes can place packets on the network segment.

The switch supports only the active paths method of determining the best return path back to a unicast source address. The active paths method looks up the best reverse path entry in the forwarding table. It does not consider alternate routes specified using routing-protocol-specific methods when determining the best return path.

If the forwarding table lists the receiving interface as the interface to use to forward the packet back to its unicast source, it is the best return path interface.

Use strict mode unicast RPF only on symmetrically routed interfaces. (For information about symmetrically routed interfaces, see [“When to Enable Unicast RPF” on page 13.](#))

For more information about strict unicast RPF, see RFC 3704, *Ingress Filtering for Multihomed Networks* at <http://www.ietf.org/rfc/rfc3704.txt>.

Unicast RPF Implementation

This section includes:

- [Unicast RPF Packet Filtering on page 13](#)
- [Bootstrap Protocol \(BOOTP\) and DHCP Requests on page 13](#)
- [Default Route Handling on page 13](#)

Unicast RPF Packet Filtering

When you enable unicast RPF on the switch, the switch handles traffic in the following manner:

- If the switch receives a packet on the interface that is the best return path to the unicast source address of that packet, the switch forwards the packet.
- If the best return path from the switch to the packet's unicast source address is not the receiving interface, the switch discards the packet.
- If the switch receives a packet that has a source IP address that does not have a routing entry in the forwarding table, the switch discards the packet.

Bootstrap Protocol (BOOTP) and DHCP Requests

Bootstrap protocol (BOOTP) and DHCP request packets are sent with a broadcast MAC address and therefore the switch does not perform unicast RPF checks on them. The switch forwards all BOOTP packets and DHCP request packets without performing unicast RPF checks.

Default Route Handling

If the best return path to the source is the default route (0.0.0.0) and the default route points to **reject**, the switch discards the packets. If the default route points to a valid network interface, the switch performs a normal unicast RPF check on the packets.

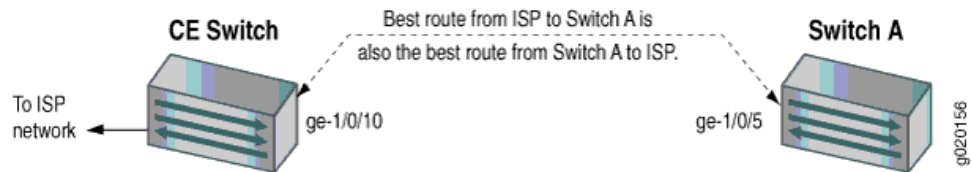
When to Enable Unicast RPF

Enable unicast RPF when you want to ensure that traffic arriving on a network interface comes from a source that resides on a network that that interface can reach. You can enable unicast RPF on untrusted interfaces to filter spoofed packets. For example, a common application for unicast RPF is to help defend an enterprise network from DoS/DDoS attacks coming from the Internet.

Enable unicast RPF only on symmetrically routed interfaces. A symmetrically routed interface uses the same route in both directions between the source and the destination, as shown in [Figure 1 on page 14](#). Symmetrical routing means that if an interface receives a packet, the switch uses the same interface to send a reply to the packet source (the

receiving interface matches the forwarding-table entry for the best return path to the source).

Figure 1: Symmetrically Routed Interfaces



Enabling unicast RPF on asymmetrically routed interfaces (where different interfaces receive a packet and reply to its source) results in packets from legitimate sources being filtered (discarded) because the best return path is not the same interface that received the packet.

The following switch interfaces are most likely to be symmetrically routed and thus are candidates for unicast RPF enabling:

- The service provider edge to a customer
- The customer edge to a service provider
- A single access point out of the network (usually on the network perimeter)
- A terminal network that has only one link



NOTE: Because unicast RPF is enabled globally on EX3200, EX4200, and EX4300 switches, ensure that *all* interfaces are symmetrically routed before you enable unicast RPF on these switches. Enabling unicast RPF on asymmetrically routed interfaces results in packets from legitimate sources being filtered.



TIP: Enabling unicast RPF as close as possible to the traffic source stops spoofed traffic before it can proliferate or reach interfaces that do not have unicast RPF enabled.

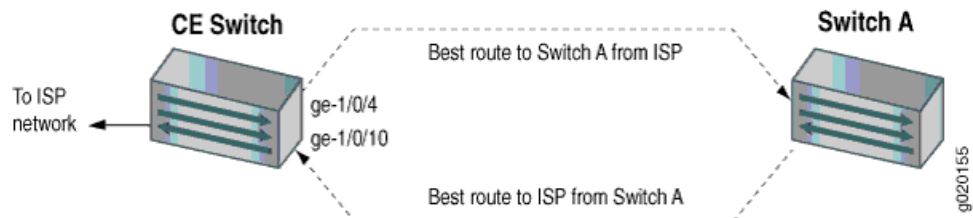
When Not to Enable Unicast RPF

Typically, you will not enable unicast RPF if:

- Switch interfaces are multihomed.
- Switch interfaces are trusted interfaces.
- BGP is carrying prefixes and some of those prefixes are not advertised or are not accepted by the ISP under its policy. (The effect in this case is the same as filtering an interface by using an incomplete access list.)
- Switch interfaces face the network core. Core-facing interfaces are usually asymmetrically routed.

An asymmetrically routed interface uses different paths to send and receive packets between the source and the destination, as shown in [Figure 2 on page 15](#). This means that if an interface receives a packet, that interface does not match the forwarding table entry as the best return path back to the source. If the receiving interface is not the best return path to the source of a packet, unicast RPF causes the switch to discard the packet even though it comes from a valid source.

Figure 2: Asymmetrically Routed Interfaces



NOTE: Do not enable unicast RPF on EX3200, EX4200, and EX4300 switches if any switch interfaces are asymmetrically routed, because unicast RPF is enabled globally on all interfaces of these switches. All switch interfaces must be symmetrically routed for you to enable unicast RPF without the risk of the switch discarding traffic that you want to forward.

Limitations of the Unicast RPF Implementation on EX3200, EX4200, and EX4300 Switches

On EX3200, EX4200, and EX4300 switches, the switch implements unicast RPF on a global basis. You cannot enable unicast RPF on a per-interface basis. Unicast RPF is globally disabled by default.

- When you enable unicast RPF on any interface, it is automatically enabled on all switch interfaces, including link aggregation groups (LAGs), integrated routing and bridging (IRB) interfaces, and routed VLAN interfaces (RVIs).
- When you disable unicast RPF on the interface (or interfaces) on which you enabled unicast RPF, it is automatically disabled on all switch interfaces.



NOTE: You must explicitly disable unicast RPF on every interface on which it was explicitly enabled or unicast RPF remains enabled on all switch interfaces.

QFX switches, OCX switches, and EX3200 and EX4200 switches do not perform unicast RPF filtering on equal-cost multipath (ECMP) traffic. The unicast RPF check examines only one best return path to the packet source, but ECMP traffic employs an address block consisting of multiple paths. Using unicast RPF to filter ECMP traffic on these switches can result in the switch discarding packets that you want to forward because the unicast RPF filter does not examine the entire ECMP address block.

- Related Documentation**
- [Example: Configuring Unicast RPF on an EX Series Switch on page 49](#)
 - [Configuring Unicast RPF \(CLI Procedure\) on page 123](#)
 - [Disabling Unicast RPF \(CLI Procedure\) on page 125](#)

Understanding IP Directed Broadcast

IP directed broadcast helps you implement remote administration tasks such as backups and wake-on-LAN (WOL) application tasks by sending broadcast packets targeted at the hosts in a specified destination subnet. IP directed broadcast packets traverse the network in the same way as unicast IP packets until they reach the destination subnet. When they reach the destination subnet and IP directed broadcast is enabled on the receiving switch, the switch translates (*explodes*) the IP directed broadcast packet into a broadcast that floods the packet on the target subnet. All hosts on the target subnet receive the IP directed broadcast packet.

This topic covers:

- [IP Directed Broadcast Overview on page 16](#)
- [IP Directed Broadcast Implementation on page 16](#)
- [When to Enable IP Directed Broadcast on page 17](#)
- [When Not to Enable IP Directed Broadcast on page 17](#)

IP Directed Broadcast Overview

IP directed broadcast packets have a destination IP address that is a valid broadcast address for the subnet that is the target of the directed broadcast (the target subnet). The intent of an IP directed broadcast is to flood the target subnet with the broadcast packets without broadcasting to the entire network. IP directed broadcast packets cannot originate from the target subnet.

When you send an IP directed broadcast packet, as it travels to the target subnet, the network forwards it in the same way as it forwards a unicast packet. When the packet reaches a switch that is directly connected to the target subnet, the switch checks to see whether IP directed broadcast is enabled on the interface that is directly connected to the target subnet:

- If IP directed broadcast is enabled on that interface, the switch broadcasts the packet on that subnet by rewriting the destination IP address as the configured broadcast IP address for the subnet. The switch converts the packet to a link-layer broadcast packet that every host on the network processes.
- If IP directed broadcast is disabled on the interface that is directly connected to the target subnet, the switch drops the packet.

IP Directed Broadcast Implementation

You configure IP directed broadcast on a per-subnet basis by enabling IP directed broadcast on the Layer 3 interface of the subnet's VLAN. When the switch that is

connected to that subnet receives a packet that has the subnet's broadcast IP address as the destination address, the switch broadcasts the packet to all hosts on the subnet.

By default, IP directed broadcast is disabled.

When to Enable IP Directed Broadcast

IP directed broadcast is disabled by default. Enable IP directed broadcast when you want to perform remote management or administration services such as backups or WOL tasks on hosts in a subnet that does not have a direct connection to the Internet.

Enabling IP directed broadcast on a subnet affects only the hosts within that subnet. Only packets received on the subnet's Layer 3 interface that have the subnet's broadcast IP address as the destination address are flooded on the subnet.

When Not to Enable IP Directed Broadcast

Typically, you do not enable IP directed broadcast on subnets that have direct connections to the Internet. Disabling IP directed broadcast on a subnet's Layer 3 interface affects only that subnet. If you disable IP directed broadcast on a subnet and a packet that has the broadcast IP address of that subnet arrives at the switch, the switch drops the broadcast packet.

If a subnet has a direct connection to the Internet, enabling IP directed broadcast on it increases the network's susceptibility to denial-of-service (DoS) attacks.

For example, a malicious attacker can spoof a source IP address (use a source IP address that is not the actual source of the transmission to deceive a network into identifying the attacker as a legitimate source) and send IP directed broadcasts containing Internet Control Message Protocol (ICMP) echo (ping) packets. When the hosts on the network with IP directed broadcast enabled receive the ICMP echo packets, they all send replies to the victim that has the spoofed source IP address. This creates a flood of ping replies in a DoS attack that can overwhelm the spoofed source address; this is known as a *smurf* attack. Another common DoS attack on exposed networks with IP directed broadcast enabled is a *fraggle* attack, which is similar to a smurf attack except that the malicious packet is a User Datagram Protocol (UDP) echo packet instead of an ICMP echo packet.

Related Documentation

- [Example: Configuring IP Directed Broadcast on a Switch on page 54](#)
- [Configuring IP Directed Broadcast \(CLI Procedure\)](#)

Understanding Interface Ranges on EX Series Switches



NOTE: This concept uses Junos OS for EX Series switches that does not support the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that supports ELS, see *Understanding Interface Ranges on EX Series Switches*. For ELS details, see *Getting Started with Enhanced Layer 2 Software*.

You can use the interface ranges to group interfaces of the same type that share a common configuration profile. This helps reduce the time and effort in configuring interfaces on Juniper Networks EX Series Ethernet switches. The configurations common to all the interfaces can be included in the interface range definition.

The interface range definition contains the name of the interface range defined, the names of the individual member interfaces that do not fall in a series of interfaces, a range of interfaces defined in the member range, and the configuration statements common to all the interfaces. An interface range defined with member ranges and individual members but without any common configurations, is also a valid definition.



NOTE: The interface range definition is supported only for Gigabit, 10-Gigabit, and Fast Ethernet interfaces.

The common configurations defined in the interface range will be overridden by the local configuration.

The defined interface ranges can be used at places where the **interface** node is used in the following configuration hierarchies:

- `ethernet-switching-options analyzer name input egress interface`
- `ethernet-switching-options analyzer name input ingress interface`
- `ethernet-switching-options analyzer output interface`
- `ethernet-switching-options bpdu-block interface`
- `ethernet-switching-options interfaces`
- `ethernet-switching-options redundant-trunk-group group-name interface`
- `ethernet-switching-options secure-access-port interface`
- `ethernet-switching-options voip interface`
- `poe interface`
- `protocols dot1x authentication interface`
- `protocols gvrp interface`
- `protocols igmp interface`
- `protocols igmp-snooping vlan vlan-name interface`
- `protocols isis interface`
- `protocols link-management peer lmp-control-channel interface`
- `protocols link-management te-link name interface`
- `protocols lldp interface`
- `protocols lldp-med interface`
- `protocols mpls interface`

- `protocols mstp interface`
- `protocols mstp msti-id interface`
- `protocols mstp msti-id vlan vlan-id interface`
- `protocols oam ethernet link-fault-management interface`
- `protocols ospf area`
- `protocols pim interface`
- `protocols rip group group-name neighbor`
- `protocols ripng group group-name neighbor`
- `protocols router-advertisement interface`
- `protocols router-discovery interface`
- `protocols rsvp interface`
- `protocols sflow interfaces`
- `protocols stp interface`
- `protocols vstp vlan vlan-id interface`
- `vlan vlan-name interface`

Related Documentation

- [Interface Ranges on page 89](#)
- [EX Series Switches Interfaces Overview on page 3](#)
- [Configuring Gigabit Ethernet Interfaces \(CLI Procedure\) on page 64](#)
- [Configuring Aggregated Ethernet Links \(CLI Procedure\) on page 109](#)
- [Configuring a Layer 3 Subinterface \(CLI Procedure\) on page 123](#)
- [interface-range on page 205](#)

802.1Q VLANs Overview

For Ethernet, Fast Ethernet, Tri-Rate Ethernet copper, Gigabit Ethernet, 10-Gigabit Ethernet, and aggregated Ethernet interfaces supporting VPLS, the Junos OS supports a subset of the IEEE 802.1Q standard for channelizing an Ethernet interface into multiple logical interfaces, allowing many hosts to be connected to the same Gigabit Ethernet switch, but preventing them from being in the same routing or bridging domain.

Related Documentation

- *Configuring Dynamic 802.1Q VLANs*
- *802.1Q VLAN IDs and Ethernet Interface Types*
- *Enabling VLAN Tagging*
- *Binding VLAN IDs to Logical Interfaces*
- *Configuring VLAN and Extended VLAN Encapsulation*

- *Guidelines for Configuring VLAN ID List-Bundled Logical Interfaces That Connect CCCs*
- *Configuring a Layer 2 VPN Routing Instance on a VLAN-Bundled Logical Interface*
- *Configuring a VLAN-Bundled Logical Interface to Support a Layer 2 VPN Routing Instance*
- *Specifying the Interface Over Which VPN Traffic Travels to the CE Router*
- *Specifying the Interface to Handle Traffic for a CCC*
- *Configuring a Layer 2 Circuit on a VLAN-Bundled Logical Interface*
- *Configuring a VLAN-Bundled Logical Interface to Support a Layer 2 VPN Routing Instance*
- *Specifying the Interface to Handle Traffic for a CCC Connected to the Layer 2 Circuit*
- *Example: Configuring a Layer 2 VPN Routing Instance on a VLAN-Bundled Logical Interface*
- *Example: Configuring a Layer 2 Circuit on a VLAN-Bundled Logical Interface*
- *Configuring Access Mode on a Logical Interface*
- *Configuring a Logical Interface for Trunk Mode*
- *Configuring the VLAN ID List for a Trunk Interface*
- *Configuring a Trunk Interface on a Bridge Network*
- *Ethernet Interfaces Feature Guide for Routing Devices*

Understanding Generic Routing Encapsulation

Generic routing encapsulation (GRE) provides a private, secure path for transporting packets through an otherwise public network by encapsulating (or tunneling) the packets.

This topic describes:

- [Overview of GRE on page 20](#)
- [GRE Tunneling on page 21](#)
- [Using a Firewall Filter to De-encapsulate GRE Traffic on a QFX5100 and OCX Series Switches on page 23](#)
- [Configuration Limitations on page 23](#)

Overview of GRE

GRE encapsulates data packets and redirects them to a device that de-encapsulates them and routes them to their final destination. This allows the source and destination switches to operate as if they have a virtual point-to-point connection with each other (because the outer header applied by GRE is transparent to the encapsulated payload packet). For example, GRE tunnels allow routing protocols such as RIP and OSPF to forward data packets from one switch to another switch across the Internet. In addition, GRE tunnels can encapsulate multicast data streams for transmission over the Internet.

GRE is described in RFC 2784 (obsoletes earlier RFCs 1701 and 1702). The switches support RFC 2784, but not completely. (For a list of limitations, see [“Configuration Limitations” on page 23.](#))

As a *tunnel source router*, the switch encapsulates a payload packet for transport through the tunnel to a destination network. The payload packet is first encapsulated in a GRE packet, and then the GRE packet is encapsulated in a delivery protocol. The switch performing the role of a *tunnel remote router* extracts the tunneled packet and forwards the packet to its destination. Note that you can use one firewall term to terminate many GRE tunnels on a QFX5100 switch.

GRE Tunneling

Data is routed by the system to the GRE endpoint over routes established in the route table. (These routes can be statically configured or dynamically learned by routing protocols such as RIP or OSPF.) When a data packet is received by the GRE endpoint, it is de-encapsulated and routed again to its destination address.

GRE tunnels are *stateless*—that is, the endpoint of the tunnel contains no information about the state or availability of the remote tunnel endpoint. Therefore, the switch operating as a tunnel source router cannot change the state of the GRE tunnel interface to down if the remote endpoint is unreachable.

For details about GRE tunneling, see:

- [Encapsulation and De-Encapsulation on the Switch on page 21](#)
- [Number of Source and Destination Tunnels Allowed on a Switch on page 21](#)
- [Class of Service on GRE Tunnels on page 22](#)
- [Applying Firewall Filters to GRE Traffic on page 22](#)

Encapsulation and De-Encapsulation on the Switch

Encapsulation—A switch operating as a tunnel source router encapsulates and forwards GRE packets as follows:

1. When a switch receives a data packet (payload) to be tunneled, it sends the packet to the tunnel interface.
2. The tunnel interface encapsulates the data in a GRE packet and adds an outer IP header.
3. The IP packet is forwarded on the basis of the destination address in the outer IP header.

De-encapsulation—A switch operating as a tunnel remote router handles GRE packets as follows:

1. When the destination switch receives the IP packet from the tunnel interface, the outer IP header and GRE header are removed.
2. The packet is routed based on the inner IP header.

Number of Source and Destination Tunnels Allowed on a Switch

QFX5100 and OCX Series switches support as many as 512 GRE tunnels, including tunnels created with a firewall filter. That is, you can create a total of 512 GRE tunnels, regardless of which method you use.

EX switches support as many as 500 GRE tunnels between switches transmitting IPv4 or IPv6 payload packets over GRE. If a passenger protocol in addition to IPv4 and IPv6 is used, you can configure up to 333 GRE tunnels between the switches.

An EX switch can have a maximum of 20 tunnel source IP addresses configured, and each tunnel source IP can be configured with up to 20 destination IP addresses on a second switch. As a result, the two connected switches can have a maximum of 400 GRE tunnels. If the first switch is also connected to a third switch, the possible maximum number of tunnels is 500.

Class of Service on GRE Tunnels

When a network experiences congestion and delay, some packets might be dropped. Junos OS class of service (CoS) divides traffic into classes to which you can apply different levels of throughput and packet loss when congestion occurs and thereby set rules for packet loss. For details about CoS, see [Junos OS CoS for EX Series Switches Overview](#).

The following CoS components are available on a switch operating as a GRE tunnel source router or GRE tunnel remote router:

- At the GRE tunnel source—On a switch operating as a tunnel source router, you can apply CoS classifiers on an *ingress port* or on a *GRE port*, with the following results on CoS component support on tunneled packets:
 - Schedulers only—Based on the CoS classification on the ingress port, you can apply CoS schedulers on a GRE port of the switch to define output queues and control the transmission of packets through the tunnel after GRE encapsulation. However, you cannot apply CoS rewrite rules to these packets.
 - Schedulers and rewrite rules—Depending on the CoS classification on the GRE port, you can apply both schedulers and rewrite rules to the encapsulated packets transmitted through the tunnel.
- At the GRE tunnel endpoint—When the switch is a tunnel remote router, you can apply CoS classifiers on the GRE port and schedulers and rewrite rules on the egress port to control the transmission of a de-encapsulated GRE packet out from the egress port.

Applying Firewall Filters to GRE Traffic

Firewall filters provide rules that define whether to permit, deny, or forward packets that are transiting an interface on a switch. (For details, see [Firewall Filters for EX Series Switches Overview](#).) Because of the encapsulation and de-encapsulation performed by GRE, you are constrained as to where you can apply a firewall filter to filter tunneled packets and which header will be affected. [Table 6 on page 22](#) identifies these constraints.

Table 6: Firewall Filter Application Points for Tunneled Packets

Endpoint Type	Ingress Interface	Egress Interface
Source (encapsulating)	inner header	outer header
Remote (de-encapsulating)	Cannot filter packets on ingress interface	inner header

Using a Firewall Filter to De-encapsulate GRE Traffic on a QFX5100 and OCX Series Switches

You can also use a firewall filter to de-encapsulate GRE traffic on switches. This feature provides significant benefits in terms of scalability, performance, and flexibility because you don't need to create a tunnel interface to perform the de-encapsulation. For example, you can terminate many tunnels from multiple source IP addresses with one firewall term. See *Configuring a Firewall Filter to De-encapsulate GRE Traffic on a QFX5100 Switch* for information about how to configure a firewall filter for this purpose.

Configuration Limitations

Table 7 on page 23 lists features that are not supported with GRE.

Table 7: Features Not Supported with GRE

EX Switches	QFX Switches
MPLS over GRE tunnels	MPLS over GRE tunnels
GRE keepalives	GRE keepalives
GRE keys, payload packet fragmentation, and sequence numbers for fragmented packets	GRE keys, payload packet fragmentation, and sequence numbers for fragmented packets
BGP dynamic tunnels	BGP dynamic tunnels
Outer IP address must be IPv4	Outer IP address must be IPv4
Virtual routing instances	
Bidirectional Forwarding Detection (BFD) protocol over GRE distributed mode	
OSPF limitation—Enabling OSPF on a GRE interface creates two equal-cost routes to the destination: one through the Ethernet network or uplink interface and the other through the tunnel interface. If data is routed through the tunnel interface, the tunnel might fail. To keep the interface operational, we recommend that you use a static route, disable OSPF on the tunnel interface, or configure the peer not to advertise the tunnel destination over the tunnel interface.	

- Related Documentation**
- [Configuring Generic Routing Encapsulation Tunneling \(CLI Procedure\) on page 130](#)
 - *Configuring Generic Routing Encapsulation Tunneling*
 - *Configuring a Firewall Filter to De-encapsulate GRE Traffic on a QFX5100 Switch*

Understanding Multicast Load Balancing on Aggregated 10-Gigabit Links for Routed Multicast Traffic on EX8200 Switches

Streaming video technology was introduced in 1997. Multicast protocols were subsequently developed to reduce data replication and network overloads. With multicasting, servers can send a single stream to a group of recipients instead of sending multiple unicast streams. While the use of streaming video technology was previously limited to occasional company presentations, multicasting has provided a boost to the technology resulting in a constant stream of movies, real-time data, news clips, and amateur videos flowing nonstop to computers, TVs, tablets, and phones. However, all of these streams quickly overwhelmed the capacity of network hardware and increased bandwidth demands leading to unacceptable blips and stutters in transmission.

To satisfy the growing bandwidth demands, multiple links were virtually aggregated to form bigger logical point-to-point link channels for the flow of data. These virtual link combinations are called multicast interfaces, also known as link aggregation groups (LAGs).

Multicast load balancing involves managing the individual links in each LAG to ensure that each link is used efficiently. Hashing algorithms continually evaluate the data stream, adjusting stream distribution over the links in the LAG, so that no link is underutilized or overutilized. Multicast load balancing is enabled by default on Juniper Networks EX8200 Ethernet Switches.

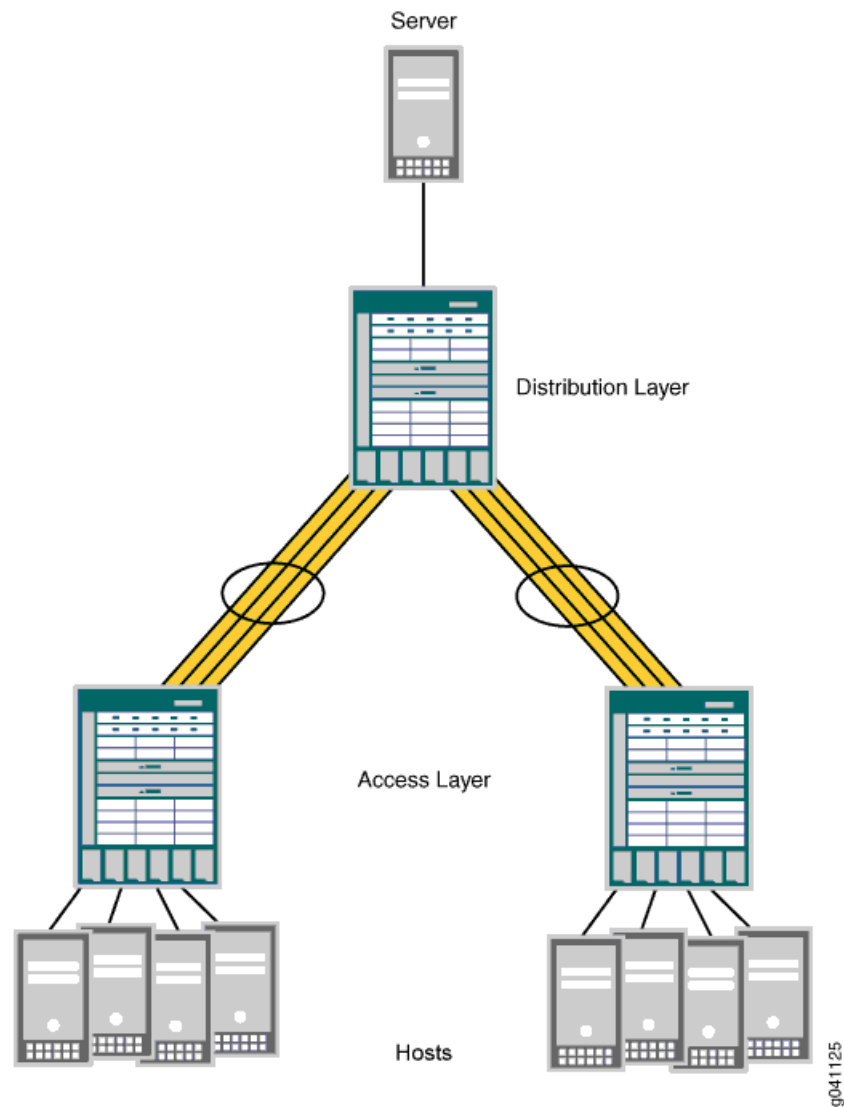
This topic includes:

- [Create LAGs for Multicasting in Increments of 10 Gigabits on page 24](#)
- [When Should I Use Multicast Load Balancing? on page 25](#)
- [How Does Multicast Load Balancing Work? on page 26](#)
- [How Do I Implement Multicast Load Balancing on an EX8200 Switch? on page 27](#)

Create LAGs for Multicasting in Increments of 10 Gigabits

The maximum link size on an EX8200 switch is 10 gigabits. If you need a larger link on an EX8200 switch, you can combine up to twelve 10-gigabit links. In the sample topology shown in [Figure 3 on page 25](#), four 10-gigabit links have been aggregated to form each 40-gigabit link.

Figure 3: 40-Gigabit LAGs on EX8200 Switches



When Should I Use Multicast Load Balancing?

Use a LAG with multicast load balancing when you need a downstream link greater than 10 gigabits. This need frequently arises when you act as a service provider or when you multicast video to a large audience.

To use multicast load balancing, you need the following:

- An EX8200 switch—Standalone switches support multicast load balancing, while Virtual Chassis does not.
- A Layer 3 routed multicast setup—For information about configuring multicasting, see [Junos OS Routing Protocols Configuration Guide](#).

- Aggregated 10-gigabit links in a LAG—For information about configuring LAGs with multicast load balancing, see [“Configuring Multicast Load Balancing for Use with Aggregated 10-Gigabit Ethernet Links on EX8200 Switches \(CLI Procedure\)”](#) on page 132.

How Does Multicast Load Balancing Work?

Juniper Networks Junos operating system (Junos OS) supports the Link Aggregation Control Protocol (LACP), which is a subcomponent of IEEE 802.3ad. LACP provides additional functionality for LAGs and is supported only on Layer 3 interfaces. When traffic can use multiple member links, traffic that is part of the same stream must always be on the same link.

Multicast load balancing uses one of seven available hashing algorithms and a technique called queue shuffling (alternating between two queues) to distribute and balance the data, directing streams over all available aggregated links. You can select one of the seven algorithms when you configure multicast load balancing, or you can use the default algorithm, `crc-sgip`, which uses a cyclic redundancy check (CRC) algorithm on the multicast packets' group IP address. We recommend that you start with the `crc-sgip` default and try other options if this algorithm does not evenly distribute the Layer 3 routed multicast traffic. Six of the algorithms are based on the hashed value of IP addresses (IPv4 or IPv6) and will produce the same result each time they are used. Only the balanced mode option produces results that vary depending on the order in which streams are added. See [Table 8 on page 26](#) for more information.

Table 8: Hashing Algorithms Used by Multicast Load Balancing

Hashing Algorithms	Based On	Best Use
<code>crc-sgip</code>	Cyclic redundancy check of multicast packets' source and group IP address	Default—high-performance management of IP traffic on 10-Gigabit Ethernet network. Predictable assignment to the same link each time. This mode is complex but yields a good distributed hash.
<code>crc-gip</code>	Cyclic redundancy check of multicast packets' group IP address	Predictable assignment to the same link each time. Try this mode when <code>crc-sgip</code> does not evenly distribute the Layer 3 routed multicast traffic and the group IP addresses vary.
<code>crc-sip</code>	Cyclic redundancy check of multicast packets' source IP address	Predictable assignment to the same link each time. Try this mode when <code>crc-sgip</code> does not evenly distribute the Layer 3 routed multicast traffic and the stream sources vary.
<code>simple-sgip</code>	XOR calculation of multicast packets' source and group IP address	Predictable assignment to the same link each time. This is a simple hashing method that might not yield as even a distribution as <code>crc-sgip</code> yields. Try this mode when <code>crc-sgip</code> does not evenly distribute the Layer 3 routed multicast traffic.
<code>simple-gip</code>	XOR calculation of multicast packets' group IP address	Predictable assignment to the same link each time. This is a simple hashing method that might not yield as even a distribution as <code>crc-gip</code> yields. Try this when <code>crc-gip</code> does not evenly distribute the Layer 3 routed multicast traffic and the group IP addresses vary.

Table 8: Hashing Algorithms Used by Multicast Load Balancing (*continued*)

Hashing Algorithms	Based On	Best Use
simple-sip	XOR calculation of multicast packets' source IP address	Predictable assignment to the same link each time. This is a simple hashing method that might not yield as even a distribution as crc-sip yields. Try this mode when crc-sip does not evenly distribute the Layer 3 routed multicast traffic and stream sources vary.
balanced	Round-robin calculation method used to identify multicast links with the least amount of traffic	Best balance is achieved, but you cannot predict which link will be consistently used because that depends on the order in which streams come online. Use when consistent assignment is not needed after every reboot.

How Do I Implement Multicast Load Balancing on an EX8200 Switch?

To implement multicast load balancing with an optimized level of throughput on an EX8200 switch, follow these recommendations:

- Allow 25 percent unused bandwidth in the aggregated link to accommodate any dynamic imbalances due to link changes caused by sharing multicast interfaces.
- For downstream links, use multicast interfaces of the same size whenever possible. Also, for downstream aggregated links, throughput is optimized when members of the aggregated link belong to the same devices.
- For upstream aggregated links, use a Layer 3 link whenever possible. Also, for upstream aggregated links, throughput is optimized when the members of the aggregated link belong to different devices.

Related Documentation

- [Example: Configuring Multicast Load Balancing for Use with Aggregated 10-Gigabit Ethernet Interfaces on EX8200 Switches on page 57](#)
- [Configuring Multicast Load Balancing for Use with Aggregated 10-Gigabit Ethernet Links on EX8200 Switches \(CLI Procedure\) on page 132](#)

Understanding How Energy Efficient Ethernet Reduces Power Consumption on Interfaces

Energy Efficient Ethernet (EEE), an Institute of Electrical and Electronics Engineers (IEEE) 802.3az standard, reduces the power consumption of physical layer devices (PHYs) during periods of low link utilization. EEE saves energy by putting part of the transmission circuit into low power mode when the link is idle.

An Ethernet link consumes power even when a link is idle. EEE provides a method to utilize power in such a way that Ethernet links use power only during data transmission. EEE specifies a signaling protocol, Low Power Idle (LPI) for achieving the power saving during the idle time of Ethernet links. EEE allows PHYs to exchange LPI indications to signal the transition to low power mode when there is no traffic. LPI indicates when a link can go idle and when the link needs to resume after a predefined delay without impacting data transmission.

The following copper PHYs are standardized by IEEE 802.3az:

- 100BASE-T
- 1000BASE-T
- 10GBASE-T

**Related
Documentation**

- [Configuring Energy Efficient Ethernet on Interfaces \(CLI Procedure\) on page 133](#)

PART 2

Configuration

- [Configuration Examples on page 31](#)
- [Configuration Tasks on page 63](#)
- [Configuration Statements on page 137](#)

CHAPTER 2

Configuration Examples

- [Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch on page 31](#)
- [Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch on page 37](#)
- [Example: Configuring Layer 3 Subinterfaces for a Distribution Switch and an Access Switch on page 42](#)
- [Example: Configuring Unicast RPF on an EX Series Switch on page 49](#)
- [Example: Configuring IP Directed Broadcast on a Switch on page 54](#)
- [Example: Configuring Multicast Load Balancing for Use with Aggregated 10-Gigabit Ethernet Interfaces on EX8200 Switches on page 57](#)

Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch

EX Series switches allow you to combine multiple Ethernet links into one logical interface for higher bandwidth and redundancy. The ports that are combined in this manner are referred to as a link aggregation group (LAG) or bundle. The number of Ethernet links you can combine into a LAG depends on your EX Series switch model. See [“Understanding Aggregated Ethernet Interfaces and LACP” on page 8](#) for more information.

This example describes how to configure uplink LAGs to connect a Virtual Chassis access switch to a Virtual Chassis distribution switch:

- [Requirements on page 31](#)
- [Overview and Topology on page 32](#)
- [Configuration on page 34](#)
- [Verification on page 36](#)
- [Troubleshooting on page 37](#)

Requirements

This example uses the following software and hardware components:

- Junos OS Release 9.0 or later for EX Series switches
- Two EX4200-48P switches
- Two EX4200-24F switches
- Four XFP uplink modules

Before you configure the LAGs, be sure you have:

- Configured the Virtual Chassis switches. See *Configuring an EX4200, EX4500, or EX4550 Virtual Chassis (CLI Procedure)*.
- Configured the uplink ports on the switches as trunk ports. See [“Configuring Gigabit Ethernet Interfaces \(CLI Procedure\)” on page 64](#).

Overview and Topology

For maximum speed and resiliency, you can combine uplinks between an access switch and a distribution switch into LAGs. Using LAGs can be particularly effective when connecting a multimember Virtual Chassis access switch to a multimember Virtual Chassis distribution switch.

The Virtual Chassis access switch in this example is composed of two member switches. Each member switch has an uplink module with two 10-Gigabit Ethernet ports. These ports are configured as trunk ports, connecting the access switch with the distribution switch.

Configuring the uplinks as LAGs has the following advantages:

- Link Aggregation Control Protocol (LACP) can optionally be configured for link negotiation.
- It doubles the speed of each uplink from 10 Gbps to 20 Gbps.
- If one physical port is lost for any reason (a cable is unplugged or a switch port fails, or one member switch is unavailable), the logical port transparently continues to function over the remaining physical port.

The topology used in this example consists of one Virtual Chassis access switch and one Virtual Chassis distribution switch. The access switch is composed of two EX4200-48P switches (SWA-0 and SWA-1), interconnected to each other with their Virtual Chassis ports (VCPs) as member switches of Host-A. The distribution switch is composed of two EX4200-24F switches (SWD-0 and SWD-1), interconnected with their VCPs as member switches of Host-D.

Each member of the access switch has an uplink module installed. Each uplink module has two ports. The uplinks are configured to act as trunk ports, connecting the access switch with the distribution switch. One uplink port from SWA-0 and one uplink port from SWA-1 are combined as LAG **ae0** to SWD-0. This link is used for one VLAN. The remaining uplink ports from SWA-0 and from SWA-1 are combined as a second LAG connection (**ae1**) to SWD-1. LAG **ae1** is used for another VLAN.



NOTE: If the remote end of the LAG link is a security device, LACP might not be supported because security devices require a deterministic configuration. In this case, do not configure LACP. All links in the LAG are permanently operational unless the switch detects a link failure within the Ethernet physical layer or data link layers.

Figure 4: Topology for LAGs Connecting an EX4200 Virtual Chassis Access Switch to an EX4200 Virtual Chassis Distribution Switch

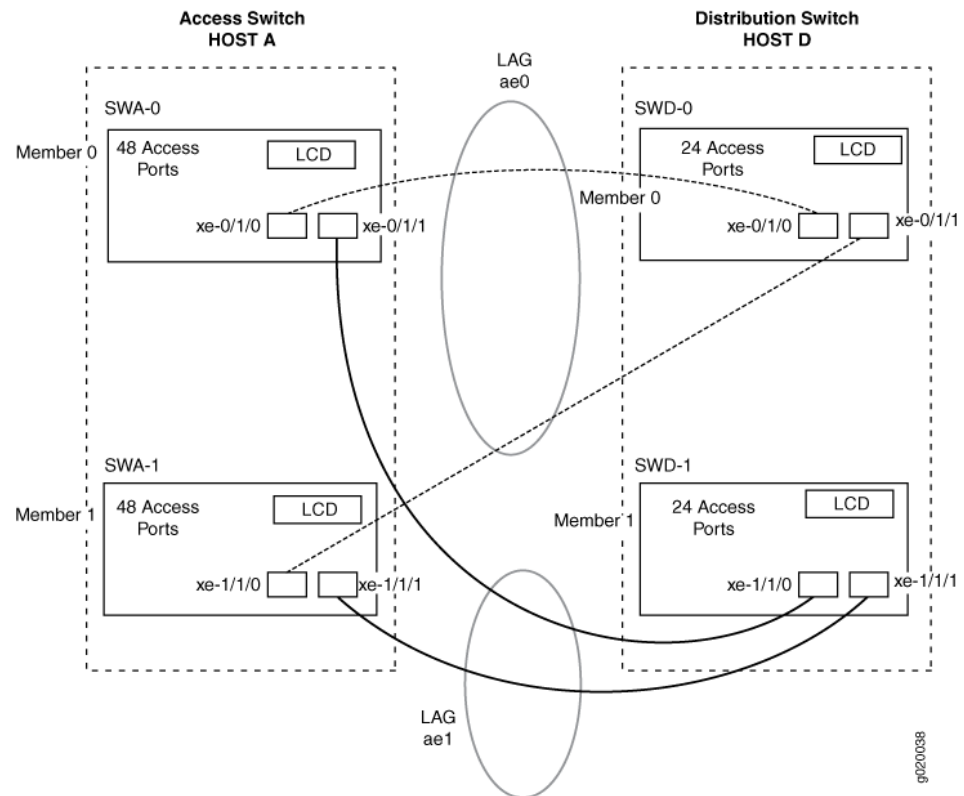


Table 9 on page 33 details the topology used in this configuration example.

Table 9: Components of the Topology for Connecting a Virtual Chassis Access Switch to a Virtual Chassis Distribution Switch

Switch	Hostname and VCID	Base Hardware	Uplink Module	Member ID	Trunk Port
SWA-0	Host-A Access switch	EX4200-48P switch	One XFP uplink module	0	xe-0/1/0 to SWD-0 xe-0/1/1 to SWD-1
	VCID 1				
SWA-1	Host-A Access switch	EX4200-48P switch	One XFP uplink module	1	xe-1/1/0 to SWD-0 xe-1/1/1 to SWD-1
	VCID 1				

Table 9: Components of the Topology for Connecting a Virtual Chassis Access Switch to a Virtual Chassis Distribution Switch (*continued*)

Switch	Hostname and VCID	Base Hardware	Uplink Module	Member ID	Trunk Port
SWD-0	Host-D Distribution switch VCID 4	EX4200 L-24F switch	One XFP uplink module	0	xe-0/1/0 to SWA-0 xe-0/1/1 to SWA-1
SWD-1	Host-D Distribution switch VCID 4	EX4200 L-24F switch	One XFP uplink module	1	xe-1/1/0 to SWA-0 xe-1/1/1 to SWA-1

Configuration

To configure two uplink LAGs from the Virtual Chassis access switch to the Virtual Chassis distribution switch:

CLI Quick Configuration

To quickly configure aggregated Ethernet high-speed uplinks between a Virtual Chassis access switch and a Virtual Chassis distribution switch, copy the following commands and paste them into the switch terminal window:

```
[edit]
set chassis aggregated-devices ethernet device-count 2
set interfaces ae0 aggregated-ether-options minimum-links 1
set interfaces ae0 aggregated-ether-options link-speed 10g
set interfaces ae1 aggregated-ether-options minimum-links 1
set interfaces ae1 aggregated-ether-options link-speed 10g
set interfaces ae0 unit 0 family inet address 192.0.2.0/25
set interfaces ae1 unit 0 family inet address 192.0.2.128/25
set interfaces xe-0/1/0 ether-options 802.3ad ae0
set interfaces xe-1/1/0 ether-options 802.3ad ae0
set interfaces xe-0/1/1 ether-options 802.3ad ae1
set interfaces xe-1/1/1 ether-options 802.3ad ae1
```

Step-by-Step Procedure

To configure aggregated Ethernet high-speed uplinks between a Virtual Chassis access switch and a Virtual Chassis distribution switch:

- Specify the number of LAGs to be created on the chassis:


```
[edit chassis]
user@Host-A# set aggregated-devices ethernet device-count 2
```
- Specify the number of links that need to be present for the **ae0** LAG interface to be up:


```
[edit interfaces]
user@Host-A# set ae0 aggregated-ether-options minimum-links 1
```
- Specify the number of links that need to be present for the **ae1** LAG interface to be up:


```
[edit interfaces]
user@Host-A# set ae1 aggregated-ether-options minimum-links 1
```
- Specify the media speed of the **ae0** link:


```
[edit interfaces]
user@Host-A# set ae0 aggregated-ether-options link-speed 10g
```

5. Specify the media speed of the **ae1** link:


```
[edit interfaces]
user@Host-A# set ae1 aggregated-ether-options link-speed 10g
```
6. Specify the interface ID of the uplinks to be included in LAG **ae0**:


```
[edit interfaces]
user@Host-A# set xe-0/1/0 ether-options 802.3ad ae0
user@Host-A# set xe-1/1/0 ether-options 802.3ad ae0
```
7. Specify the interface ID of the uplinks to be included in LAG **ae1**:


```
[edit interfaces]
user@Host-A# set xe-0/1/1 ether-options 802.3ad ae1
user@Host-A# set xe-1/1/1 ether-options 802.3ad ae1
```
8. Specify that LAG **ae0** belongs to the subnet for the employee broadcast domain:


```
[edit interfaces]
user@Host-A# set ae0 unit 0 family inet address 192.0.2.0/25
```
9. Specify that LAG **ae1** belongs to the subnet for the guest broadcast domain:


```
[edit interfaces]
user@Host-A# set ae1 unit 0 family inet address 192.0.2.128/25
```

Results

Display the results of the configuration:

```
[edit]
chassis {
  aggregated-devices {
    ethernet {
      device-count 2;
    }
  }
}
interfaces {
  ae0 {
    aggregated-ether-options {
      link-speed 10g;
      minimum-links 1;
    }
    unit 0 {
      family inet {
        address 192.0.2.0/25;
      }
    }
  }
  ae1 {
    aggregated-ether-options {
      link-speed 10g;
      minimum-links 1;
    }
    unit 0 {
      family inet {
        address 192.0.2.128/25;
      }
    }
  }
  xe-0/1/0 {
```

```

        ether-options {
            802.3ad ae0;
        }
    }
    xe-1/1/0 {
        ether-options {
            802.3ad ae0;
        }
    }
    xe-0/1/1 {
        ether-options {
            802.3ad ae1;
        }
    }
    xe-1/1/1 {
        ether-options {
            802.3ad ae1;
        }
    }
}

```

Verification

To verify that switching is operational and two LAGs have been created, perform these tasks:

- [Verifying That LAG ae0 Has Been Created on page 36](#)
- [Verifying That LAG ae1 Has Been Created on page 36](#)

Verifying That LAG ae0 Has Been Created

Purpose Verify that LAG **ae0** has been created on the switch.

Action `show interfaces ae0 terse`

Interface	Admin	Link	Proto	Local	Remote
ae0	up	up			
ae0.0	up	up	inet	192.0.2.0/25	

Meaning The output confirms that the **ae0** link is up and shows the **family** and IP address assigned to this link.

Verifying That LAG ae1 Has Been Created

Purpose Verify that LAG **ae1** has been created on the switch

Action `show interfaces ae1 terse`

Interface	Admin	Link	Proto	Local	Remote
ae1	up	down			
ae1.0	up	down	inet	192.0.2.128/25	

Meaning The output shows that the **ae1** link is down.

Troubleshooting

Troubleshooting a LAG That Is Down

Problem	The <code>show interfaces terse</code> command shows that the LAG is down .
Solution	<p>Check the following:</p> <ul style="list-style-type: none"> • Verify that there is no configuration mismatch. • Verify that all member ports are up. • Verify that a LAG is part of family ethernet switching (Layer 2 LAG) or family inet (Layer 3 LAG). • Verify that the LAG member is connected to the correct LAG at the other end. • Verify that the LAG members belong to the same switch (or the same Virtual Chassis).
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring an EX4200 Virtual Chassis with a Master and Backup in a Single Wiring Closet</i> • Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch on page 37 • <i>Example: Connecting an Access Switch to a Distribution Switch.</i> • <i>Virtual Chassis Cabling Configuration Examples for EX4200 Switches</i> • <i>Installing an Uplink Module in an EX4200 Switch</i> • <i>Uplink Modules in EX4200 Switches</i>

Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch

EX Series switches allow you to combine multiple Ethernet links into one logical interface for higher bandwidth and redundancy. The ports that are combined in this manner are referred to as a link aggregation group (LAG) or bundle. EX Series switches allow you to further enhance these links by configuring Link Aggregation Control Protocol (LACP).

This example describes how to overlay LACP on the LAG configurations that were created in “[Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch](#)” on [page 31](#):

- [Requirements on page 38](#)
- [Overview and Topology on page 38](#)
- [Configuring LACP for the LAGs on the Virtual Chassis Access Switch on page 38](#)
- [Configuring LACP for the LAGs on the Virtual Chassis Distribution Switch on page 39](#)

- [Verification on page 40](#)
- [Troubleshooting on page 41](#)

Requirements

This example uses the following software and hardware components:

- Junos OS Release 9.0 or later for EX Series switches
- Two EX4200-48P switches
- Two EX4200-24F switches
- Four EX Series XFP uplink modules

Before you configure LACP, be sure you have:

- Set up the Virtual Chassis switches. See *Configuring an EX4200, EX4500, or EX4550 Virtual Chassis (CLI Procedure)*.
- Configured the uplink ports on the switches as trunk ports. See *“Configuring Gigabit Ethernet Interfaces (CLI Procedure)” on page 64*.
- Configured the LAGs. See *“Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch” on page 31*.

Overview and Topology

This example assumes that you are familiar with *“Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch” on page 31*. The topology in this example is exactly the same as the topology in that other example. This example shows how to use LACP to enhance the LAG functionality.

LACP exchanges are made between *actors* (the transmitting link) and *partners* (the receiving link). The LACP mode can be either active or passive.



.....

NOTE: If the actor and partner are both in passive mode, they do not exchange LACP packets, which results in the aggregated Ethernet links not coming up. By default, LACP is in passive mode. To initiate transmission of LACP packets and responses to LACP packets, you must enable LACP in active mode.

.....

By default, the actor and partner send LACP packets every second.

The interval can be fast (every second) or slow (every 30 seconds).

Configuring LACP for the LAGs on the Virtual Chassis Access Switch

To configure LACP for the access switch LAGs, perform these tasks:

CLI Quick Configuration To quickly configure LACP for the access switch LAGs, copy the following commands and paste them into the switch terminal window:

```
[edit]
set interfaces ae0 aggregated-ether-options lacp active periodic fast
set interfaces ae1 aggregated-ether-options lacp active periodic fast
```

Step-by-Step Procedure To configure LACP for Host-A LAGs ae0 and ae1:

1. Specify the aggregated Ethernet options for both bundles:

```
[edit interfaces]
user@Host-A#set ae0 aggregated-ether-options lacp active periodic fast
user@Host-A#set ae1 aggregated-ether-options lacp active periodic fast
```

Results Display the results of the configuration:

```
[edit interfaces]
user@Host-A# show
ae0 {
  aggregated-ether-options {
    lacp {
      active;
      periodic fast;
    }
  }
}
ae1 {
  aggregated-ether-options {
    lacp {
      active;
      periodic fast;
    }
  }
}
```

Configuring LACP for the LAGs on the Virtual Chassis Distribution Switch

To configure LACP for the two uplink LAGs from the Virtual Chassis access switch to the Virtual Chassis distribution switch, perform these tasks:

CLI Quick Configuration To quickly configure LACP for the distribution switch LAGs, copy the following commands and paste them into the switch terminal window:

```
[edit interfaces]
set ae0 aggregated-ether-options lacp passive periodic fast
set ae1 aggregated-ether-options lacp passive periodic fast
```

Step-by-Step Procedure To configure LACP for Host D LAGs ae0 and ae1:

1. Specify the aggregated Ethernet options for both bundles:

```
[edit interfaces]
user@Host-D#set ae0 aggregated-ether-options lacp passive periodic fast
user@Host-D#set ae1 aggregated-ether-options lacp passive periodic fast
```

Results Display the results of the configuration:

```
[edit interfaces]
```

```

user@Host-D# show
ae0 {
  aggregated-ether-options {
    lACP {
      passive;
      periodic fast;
    }
  }
}
ae1 {
  aggregated-ether-options {
    lACP {
      passive
      periodic fast;
    }
  }
}

```

Verification

To verify that LACP packets are being exchanged, perform these tasks:

- [Verifying the LACP Settings on page 40](#)
- [Verifying That the LACP Packets Are Being Exchanged on page 40](#)

Verifying the LACP Settings

Purpose Verify that LACP has been set up correctly.

Action Use the **show lACP interfaces *interface-name*** command to check that LACP has been enabled as active on one end.

```
user@Host-A> show lACP interfaces xe-0/1/0
```

Aggregated interface: ae0

LACP state:	Role	Exp	Def	Dist	Col	Syn	Aggr	Timeout	Activity
xe-0/1/0	Actor	No	Yes	No	No	No	Yes	Fast	Active
xe-0/1/0	Partner	No	Yes	No	No	No	Yes	Fast	Passive
LACP protocol:	Receive State		Transmit State		Mux State				
xe-0/1/0	Defaulted		Fast periodic		Detached				

Meaning The output indicates that LACP has been set up correctly and is active at one end.

Verifying That the LACP Packets Are Being Exchanged

Purpose Verify that LACP packets are being exchanged.

Action Use the **show interfaces aex statistics** command to display LACP information.

```
user@Host-A> show interfaces ae0 statistics
```

```
Physical interface: ae0, Enabled, Physical link is Down
Interface index: 153, SNMP ifIndex: 30
Link-level type: Ethernet, MTU: 1514, Speed: Unspecified, Loopback: Disabled,
Source filtering: Disabled, Flow control: Disabled, Minimum links needed: 1,
Minimum bandwidth needed: 0
Device flags   : Present Running
Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
Current address: 02:19:e2:50:45:e0, Hardware address: 02:19:e2:50:45:e0
Last flapped   : Never
Statistics last cleared: Never
  Input packets : 0
  Output packets: 0
Input errors: 0, Output errors: 0

Logical interface ae0.0 (Index 71) (SNMP ifIndex 34)
Flags: Hardware-Down Device-Down SNMP-Traps Encapsulation: ENET2
Statistics      Packets      pps      Bytes      bps
Bundle:
  Input :           0           0           0           0
  Output:           0           0           0           0
Protocol inet
Flags: None
Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
Destination: 10.10.10/24, Local: 10.10.10.1, Broadcast: 10.10.10.255
```

Meaning The output here shows that the link is down and that no protocol data units (PDUs) are being exchanged.

Troubleshooting

To troubleshoot a nonworking LACP link, perform these tasks:

- [Troubleshooting a Nonworking LACP Link on page 41](#)

[Troubleshooting a Nonworking LACP Link](#)

Problem The LACP link is not working.

Solution Check the following:

- Remove the LACP configuration and verify whether the static LAG is up.
- Verify that LACP is configured at both ends.
- Verify that LACP is not passive at both ends.
- Verify whether LACP protocol data units (PDUs) are being exchanged by running the **monitor traffic-interface lag-member detail** command.

Related Documentation

- *Example: Connecting an Access Switch to a Distribution Switch*
- *Virtual Chassis Cabling Configuration Examples for EX4200 Switches*

- [Installing an Uplink Module in an EX4200 Switch](#)
- [Understanding Aggregated Ethernet Interfaces and LACP on page 8](#)

Example: Configuring Layer 3 Subinterfaces for a Distribution Switch and an Access Switch

In a large LAN, you commonly need to partition the network into multiple VLANs. You can configure Layer 3 subinterfaces to route traffic between the VLANs. In one common topology, known as a “router on a stick” or a “one-armed router,” you connect a router to an access switch with connections to multiple VLANs.

This example describes how to create Layer 3 subinterfaces on trunk interfaces of a distribution switch and access switch so that you can route traffic among multiple VLANs:

- [Requirements on page 42](#)
- [Overview and Topology on page 42](#)
- [Configuring the Access Switch Subinterfaces on page 43](#)
- [Configuring the Distribution Switch Subinterfaces on page 45](#)
- [Verification on page 47](#)

Requirements

This example uses the following hardware and software components:

- For the distribution switch, one EX4200-24F switch. This model is designed to be used as a distribution switch for aggregation or collapsed core network topologies and in space-constrained data centers. It has twenty-four 1-Gigabit Ethernet fiber SFP ports and an EX-UM-2XFP uplink module with two 10-Gigabit Ethernet XFP ports.
- For the access switch, any Layer 2 switch that supports 802.1Q VLAN tags.
- Junos OS Release 9.2 or later for EX Series switches.

Before you connect the switches, make sure you have:

- Connected the two switches.
- Configured the necessary VLANs. See *Configuring VLANs for EX Series Switches (CLI Procedure)* or *Configuring VLANs for EX Series Switches (J-Web Procedure)*.

Overview and Topology

In a large office with multiple buildings and VLANs, you commonly aggregate traffic from a number of access switches into a distribution switch. This configuration example shows a simple topology to illustrate how to connect a single Layer 2 access switch connected to multiple VLANs to a distribution switch, enabling traffic to pass between those VLANs.

In the example topology, the LAN is segmented into five VLANs, all associated with interfaces on the access switch. One 1-Gigabit Ethernet port on the access switch's uplink module connects to one 1-Gigabit Ethernet port on the distribution switch.

Table 10 on page 43 lists the settings for the example topology.

Table 10: Components of the Topology for Creating Layer 3 Subinterfaces on an Access Switch and a Distribution Switch

Property	Settings
Access switch hardware	Any Layer 2 switch with multiple 1-Gigabit Ethernet ports and at least one 1-Gigabit Ethernet uplink module
Distribution switch hardware	EX4200-24F, 24 1-Gigabit Ethernet fiber SPF ports (ge-0/0/0 through ge-0/0/23); one 2-port 10-Gigabit Ethernet XFP uplink module (EX-UM-4SFP)
VLAN names and tag IDs	vlan1, tag 101 vlan2, tag 102 vlan3, tag 103 vlan4, tag 104 vlan5, tag 105
VLAN subnets	vlan1: 1.1.1.0/24 (addresses 1.1.1.1 through 1.1.1.254) vlan2: 2.1.1.0/24 (addresses 2.1.1.1 through 2.1.1.254) vlan3: 3.1.1.0/24 (addresses 3.1.1.1 through 3.1.1.254) vlan4: 4.1.1.0/24 (addresses 4.1.1.1 through 4.1.1.254) vlan5: 5.1.1.0/24 (addresses 5.1.1.1 through 5.1.1.254)
Port interfaces	On the access switch: ge-0/1/0 On the distribution switch: ge-0/0/0

Configuring the Access Switch Subinterfaces

CLI Quick Configuration To quickly create and configure subinterfaces on the access switch, copy the following commands and paste them into the switch terminal window:

```
[edit]
set interfaces ge-0/1/0 vlan-tagging
set interfaces ge-0/1/0 unit 0 vlan-id 101 family inet address 1.1.1.1/24
set interfaces ge-0/1/0 unit 1 vlan-id 102 family inet address 2.1.1.1/24
set interfaces ge-0/1/0 unit 2 vlan-id 103 family inet address 3.1.1.1/24
set interfaces ge-0/1/0 unit 3 vlan-id 104 family inet address 4.1.1.1/24
set interfaces ge-0/1/0 unit 4 vlan-id 105 family inet address 5.1.1.1/24
```

Step-by-Step Procedure To configure the subinterfaces on the access switch:

- On the trunk interface of the access switch, enable VLAN tagging:


```
[edit interfaces ge-0/1/0]
user@access-switch# set vlan-tagging
```
- Bind vlan1's VLAN ID to the logical interface:


```
[edit interfaces ge-0/1/0]
user@access-switch# set unit 0 vlan-id 101
```

3. Set vlan1's subinterface IP address:

```
[edit interfaces ge-0/1/0]
user@access-switch# set unit 0 family inet address 1.1.1.1/24
```
4. Bind vlan2's VLAN ID to the logical interface:

```
[edit interfaces ge-0/1/0]
user@access-switch# set unit 1 vlan-id 102
```
5. Set vlan2's subinterface IP address:

```
[edit interfaces ge-0/1/0]
user@access-switch# set unit 1 family inet address 2.1.1.1/24
```
6. Bind vlan3's VLAN ID to the logical interface:

```
[edit interfaces ge-0/1/0]
user@access-switch# set unit 2 vlan-id 103
```
7. Set vlan3's subinterface IP address:

```
[edit interfaces ge-0/1/0]
user@access-switch# set unit 2 family inet address 3.1.1.1/24
```
8. Bind vlan4's VLAN ID to the logical interface:

```
[edit interfaces ge-0/1/0]
user@access-switch# set unit 3 vlan-id 104
```
9. Set vlan4's subinterface IP address:

```
[edit interfaces ge-0/1/0]
user@access-switch# set unit 3 family inet address 4.1.1.1/24
```
10. Bind vlan5's VLAN ID to the logical interface:

```
[edit interfaces ge-0/1/0]
user@access-switch# set unit 4 vlan-id 105
```
11. Set vlan5's subinterface IP address:

```
[edit interfaces ge-0/1/0]
user@access-switch# set unit 4 family inet address 5.1.1.1/24
```

Results

Check the results of the configuration:

```
user@access-switch> show configuration
interfaces {
  ge-0/1/0 {
    vlan-tagging;
    unit 0 {
      vlan-id 101;
      family inet {
        address 1.1.1.1/24;
      }
    }
    unit 1 {
      vlan-id 102;
      family inet {
        address 2.1.1.1/24;
      }
    }
    unit 2 {
      vlan-id 103;
      family inet {
```

```

        address 3.1.1.1/24;
    }
}
unit 3 {
    vlan-id 104;
    family inet {
        address 4.1.1.1/24;
    }
}
unit 4 {
    vlan-id 105;
    family inet {
        address 5.1.1.1/24;
    }
}
}

```

Configuring the Distribution Switch Subinterfaces

CLI Quick Configuration To quickly create and configure subinterfaces on the distribution switch, copy the following commands and paste them into the switch terminal window:

```

[edit]
set interfaces ge-0/0/0 vlan-tagging
set interfaces ge-0/0/0 unit 0 vlan-id 101 family inet address 1.1.1.2/24
set interfaces ge-0/0/0 unit 1 vlan-id 102 family inet address 2.1.1.2/24
set interfaces ge-0/0/0 unit 2 vlan-id 103 family inet address 3.1.1.2/24
set interfaces ge-0/0/0 unit 3 vlan-id 104 family inet address 4.1.1.2/24
set interfaces ge-0/0/0 unit 4 vlan-id 105 family inet address 5.1.1.2/24

```

Step-by-Step Procedure To configure subinterfaces on the distribution switch:

1. On the trunk interface of the distribution switch, enable VLAN tagging:

```

[edit interfaces ge-0/0/0]
user@distribution-switch# set vlan-tagging

```

2. Bind vlan1's VLAN ID to the logical interface:

```

[edit interfaces ge-0/0/0]
user@distribution-switch# set unit 0 vlan-id 101

```

3. Set vlan1's subinterface IP address:

```

[edit interfaces ge-0/0/0]
user@distribution-switch# set unit 0 family inet address 1.1.1.2/24

```

4. Bind vlan2's VLAN ID to the logical interface:

```

[edit interfaces ge-0/0/0]
user@distribution-switch# set unit 1 vlan-id 102

```

5. Set vlan2's subinterface IP address:

```

[edit interfaces ge-0/0/0]
user@distribution-switch# set unit 1 family inet address 2.1.1.2/24

```

6. Bind vlan3's VLAN ID to the logical interface:

```

[edit interfaces ge-0/0/0]
user@distribution-switch# set unit 2 vlan-id 103

```

7. Set vlan3's subinterface IP address:

```

[edit interfaces ge-0/0/0]
user@distribution-switch# set unit 2 family inet address 3.1.1.2/24

```

8. Bind vlan4's VLAN ID to the logical interface:

```
[edit interfaces ge-0/0/0]  
user@distribution-switch# set unit 3 vlan-id 104
```
9. Set vlan4's subinterface IP address:

```
[edit interfaces ge-0/0/0]  
user@distribution-switch# set unit 3 family inet address 4.1.1.2/24
```
10. Bind vlan5's VLAN ID to the logical interface:

```
[edit interfaces ge-0/0/0]  
user@distribution-switch# set unit 4 vlan-id 105
```
11. Set vlan5's subinterface IP address:

```
[edit interfaces ge-0/0/0]  
user@distribution-switch# set unit 4 family inet address 5.1.1.2/24
```

Results

```
user@distribution-switch> show configuration
```

```
interfaces {  
  ge-0/0/0 {  
    vlan-tagging;  
    unit 0 {  
      vlan-id 101;  
      family inet {  
        address 1.1.1.2/24;  
      }  
    }  
    unit 1 {  
      vlan-id 102;  
      family inet {  
        address 2.1.1.2/24;  
      }  
    }  
    unit 2 {  
      vlan-id 103;  
      family inet {  
        address 3.1.1.2/24;  
      }  
    }  
    unit 3 {  
      vlan-id 104;  
      family inet {  
        address 4.1.1.2/24;  
      }  
    }  
    unit 4 {  
      vlan-id 105;  
      family inet {  
        address 5.1.1.2/24;  
      }  
    }  
  }  
}
```


Action Ping from the access switch to the distribution switch on each subinterface.

1. From the access switch, ping the address of the vlan1 subinterface on the distribution switch:

```
user@access-switch> ping 1.1.1.2 count 4
PING 1.1.1.2 (1.1.1.2): 56 data bytes
64 bytes from 1.1.1.2: icmp_seq=0 ttl=64 time=0.333 ms
64 bytes from 1.1.1.2: icmp_seq=1 ttl=64 time=0.113 ms
64 bytes from 1.1.1.2: icmp_seq=2 ttl=64 time=0.112 ms
64 bytes from 1.1.1.2: icmp_seq=3 ttl=64 time=0.158 ms

--- 1.1.1.2 ping statistics ---
4 packets transmitted, 4 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.112/0.179/0.333/0.091 ms
```

2. From the access switch, ping the address of the vlan2 subinterface on the distribution switch:

```
user@access-switch> ping 2.1.1.2 count 4
PING 2.1.1.2 (2.1.1.2): 56 data bytes
64 bytes from 2.1.1.2: icmp_seq=0 ttl=64 time=0.241 ms
64 bytes from 2.1.1.2: icmp_seq=1 ttl=64 time=0.113 ms
64 bytes from 2.1.1.2: icmp_seq=2 ttl=64 time=0.162 ms
64 bytes from 2.1.1.2: icmp_seq=3 ttl=64 time=0.167 ms

--- 2.1.1.2 ping statistics ---
4 packets transmitted, 4 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.113/0.171/0.241/0.046 ms
```

3. From the access switch, ping the address of the vlan3 subinterface on the distribution switch:

```
user@access-switch> ping 3.1.1.2 count 4
PING 3.1.1.2 (3.1.1.2): 56 data bytes
64 bytes from 3.1.1.2: icmp_seq=0 ttl=64 time=0.341 ms
64 bytes from 3.1.1.2: icmp_seq=1 ttl=64 time=0.162 ms
64 bytes from 3.1.1.2: icmp_seq=2 ttl=64 time=0.112 ms
64 bytes from 3.1.1.2: icmp_seq=3 ttl=64 time=0.208 ms

--- 3.1.1.2 ping statistics ---
4 packets transmitted, 4 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.112/0.206/0.341/0.085 ms
```

4. From the access switch, ping the address of the vlan4 subinterface on the distribution switch:

```
user@access-switch> ping 4.1.1.2 count 4
PING 4.1.1.2 (4.1.1.2): 56 data bytes
64 bytes from 4.1.1.2: icmp_seq=0 ttl=64 time=0.226 ms
64 bytes from 4.1.1.2: icmp_seq=1 ttl=64 time=0.166 ms
64 bytes from 4.1.1.2: icmp_seq=2 ttl=64 time=0.107 ms
64 bytes from 4.1.1.2: icmp_seq=3 ttl=64 time=0.221 ms

--- 4.1.1.2 ping statistics ---
4 packets transmitted, 4 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.107/0.180/0.226/0.048 ms
```

5. From the access switch, ping the address of the vlan5 subinterface on the distribution switch:

```
user@access-switch> ping 5.1.1.2 count 4
```

```

PING 5.1.1.2 (5.1.1.2): 56 data bytes
64 bytes from 5.1.1.2: icmp_seq=0 ttl=64 time=0.224 ms
64 bytes from 5.1.1.2: icmp_seq=1 ttl=64 time=0.104 ms
64 bytes from 5.1.1.2: icmp_seq=2 ttl=64 time=0.102 ms
64 bytes from 5.1.1.2: icmp_seq=3 ttl=64 time=0.170 ms

--- 5.1.1.2 ping statistics ---
4 packets transmitted, 4 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.102/0.150/0.224/0.051 ms

```

Meaning If all the ping packets are transmitted and are received by the destination address, the subinterfaces are up and working.

- Related Documentation**
- [Understanding Layer 3 Logical Interfaces](#)
 - [Example: Connecting an Access Switch to a Distribution Switch](#)
 - [Configuring a Layer 3 Logical Interface](#)
 - [Configuring a Layer 3 Subinterface \(CLI Procedure\) on page 123](#)

Example: Configuring Unicast RPF on an EX Series Switch

Unicast reverse-path forwarding (RPF) helps protect the switch against denial-of-service (DoS) and distributed denial-of-service (DDoS) attacks by verifying the unicast source address of each packet that arrives on an ingress interface where unicast RPF is enabled.

This example shows how to help defend the switch ingress interfaces against denial-of-service (DoS) and distributed denial-of-service (DDoS) attacks by configuring unicast reverse-path forwarding (RPF) on a customer-edge interface to filter incoming traffic:

- [Requirements on page 49](#)
- [Overview and Topology on page 50](#)
- [Configuration on page 50](#)
- [Verification on page 51](#)

Requirements

This example uses the following software and hardware components:

- Junos OS Release 10.1 or later for EX Series switches
- Two EX8200 switches

Before you begin, be sure you have:

- Connected the two switches by symmetrically routed interfaces.
- Ensured that the interface on which you will configure unicast RPF is symmetrically routed.

Overview and Topology

Large amounts of unauthorized traffic such as attempts to flood a network with fake (bogus) service requests in a denial-of-service (DoS) attack can consume network resources and deny service to legitimate users. One way to help prevent DoS and distributed denial-of-service (DDoS) attacks is to verify that incoming traffic originates from legitimate network sources.

Unicast RPF helps ensure that a traffic source is legitimate (authorized) by comparing the source address of each packet that arrives on an interface to the forwarding-table entry for its source address. If the switch uses the same interface that the packet arrived on to reply to the packet's source, this verifies that the packet originated from an authorized source, and the switch forwards the packet. If the switch does not use the same interface that the packet arrived on to reply to the packet's source, the packet might have originated from an unauthorized source, and the switch discards the packet.

This example uses two EX8200 switches. On EX3200 and EX4200 switches, you cannot configure individual interfaces for unicast RPF. On EX3200 and EX4200 switches, the switch applies unicast RPF globally to all interfaces on the switch. See [“Understanding Unicast RPF” on page 12](#) for more information on limitations regarding the configuration of unicast RPF on EX3200 and EX4200 switches.

In this example, an enterprise network's system administrator wants to protect Switch A against potential DoS and DDoS attacks from the Internet. The administrator configures unicast RPF on interface **ge-1/0/10** on Switch A. Packets arriving on interface **ge-1/0/10** on Switch A from the Switch B source also use incoming interface **ge-1/0/10** as the best return path to send packets back to the source.

The topology of this configuration example uses two EX8200 switches, Switch A and Switch B, connected by symmetrically routed interfaces:

- Switch A is on the edge of an enterprise network. The interface **ge-1/0/10** on Switch A connects to the interface **ge-1/0/5** on Switch B.
- Switch B is on the edge of the service provider network that connects the enterprise network to the Internet.

Configuration

To enable unicast RPF, perform these tasks:

CLI Quick Configuration

To quickly configure unicast RPF on Switch A, copy the following command and paste it into the switch terminal window:

```
[edit interfaces]  
set ge-1/0/10 unit 0 family inet rpf-check
```

Step-by-Step To configure unicast RPF on Switch A:

Procedure

1. Enable unicast RPF on interface **ge-1/0/10**:

```
[edit interfaces]
user@switch# set ge-1/0/10 unit 0 family inet rpf-check
```

Results Check the results:

```
[edit interfaces]
user@switch# show
ge-1/0/10 {
  unit 0 {
    family inet {
      rpf-check;
    }
  }
}
```

Verification

To confirm that the configuration is correct, perform these tasks:

- [Verifying That Unicast RPF Is Enabled on the Switch on page 51](#)

[Verifying That Unicast RPF Is Enabled on the Switch](#)

Purpose Verify that unicast RPF is enabled.

Action Verify that unicast RPF is enabled on interface **ge-1/0/10** by using the **show interfaces ge-1/0/10 extensive** or **show interfaces ge-1/0/10 detail** command.

```

user@switch> show interfaces ge-1/0/10 extensive
Physical interface: ge-1/0/10, Enabled, Physical link is Down
  Interface index: 139, SNMP ifIndex: 58, Generation: 140
  Link-level type: Ethernet, MTU: 1514, Speed: Auto, MAC-REWRITE Error: None,
  Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled,
  Auto-negotiation: Enabled, Remote fault: Online
  Device flags   : Present Running
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues    : 8 supported, 8 maximum usable queues
  Hold-times    : Up 0 ms, Down 0 ms
  Current address: 00:19:e2:50:95:ab, Hardware address: 00:19:e2:50:95:ab
  Last flapped  : Never
  Statistics last cleared: Never
  Traffic statistics:
    Input bytes :                0                0 bps
    Output bytes :                0                0 bps
    Input packets:                0                0 pps
    Output packets:                0                0 pps
  IPv6 transit statistics:
    Input bytes :                0
    Output bytes :                0
    Input packets:                0
    Output packets:                0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
    L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
    FIFO errors: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,

    FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
  Egress queues: 8 supported, 4 in use
  Queue counters:      Queued packets  Transmitted packets      Dropped packets

    0 best-effort                0                0                0
    1 assured-forw                0                0                0
    5 expedited-fo                0                0                0
    7 network-cont                0                0                0

  Active alarms : LINK
  Active defects : LINK
  MAC statistics:
    Receive          Transmit
    Total octets      0                0
    Total packets     0                0
    Unicast packets   0                0
    Broadcast packets 0                0
    Multicast packets 0                0
    CRC/Align errors  0                0
    FIFO errors       0                0
    MAC control frames 0                0
    MAC pause frames   0                0
    Oversized frames   0
    Jabber frames      0

```

```

Fragment frames                                0
VLAN tagged frames                             0
Code violations                                0
Filter statistics:
Input packet count                             0
Input packet rejects                           0
Input DA rejects                               0
Input SA rejects                               0
Output packet count                            0
Output packet pad count                        0
Output packet error count                      0
CAM destination filters: 0, CAM source filters: 0
Autonegotiation information:
Negotiation status: Incomplete
Packet Forwarding Engine configuration:
Destination slot: 1

Logical interface ge-1/0/10.0 (Index 69) (SNMP ifIndex 59) (Generation 135)
Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
Traffic statistics:
Input bytes :                                0
Output bytes :                               0
Input packets:                              0
Output packets:                              0
IPv6 transit statistics:
Input bytes :                                0
Output bytes :                               0
Input packets:                              0
Output packets:                              0
Local statistics:
Input bytes :                                0
Output bytes :                               0
Input packets:                              0
Output packets:                              0
Transit statistics:
Input bytes :                                0
Output bytes :                               0
Input packets:                              0
Output packets:                              0
IPv6 transit statistics:
Input bytes :                                0
Output bytes :                               0
Input packets:                              0
Output packets:                              0
Protocol inet, Generation: 144, Route table: 0
Flags: uRPF
Addresses, Flags: Is-Preferred Is-Primary

```

Meaning The second-to-last line of the display shows the unicast RPF flag enabled, confirming that unicast RPF is enabled on interface **ge-1/0/10**.

Related Documentation

- [Configuring Unicast RPF \(CLI Procedure\) on page 123](#)
- [Disabling Unicast RPF \(CLI Procedure\) on page 125](#)

Example: Configuring IP Directed Broadcast on a Switch

IP directed broadcast provides a method of sending broadcast packets to hosts on a specified subnet without broadcasting those packets to hosts on the entire network.

This example shows how to enable a subnet to receive IP directed broadcast packets so you can perform backups and other network management tasks remotely:

- [Requirements on page 54](#)
- [Overview and Topology on page 54](#)
- [Configuration on page 55](#)

Requirements

This example uses the following software and hardware components:

- Junos OS Release 9.4 or later for EX Series switches or Junos OS Release 15.1X53-D10 for QFX10000 switches.
- One PC
- One EX Series switch or QFX10000 switch

Before you configure IP directed broadcast for a subnet:

- Ensure that the subnet does not have a direct connection to the Internet.
- Configure routed VLAN interfaces (RVIs) for the ingress and egress VLANs on the switch. See *Configuring Routed VLAN Interfaces (CLI Procedure)* or *Configuring VLANs for EX Series Switches (J-Web Procedure)*.

Overview and Topology

You might want to perform remote administration tasks such as backups and wake-on-LAN (WOL) application tasks to manage groups of clients on a subnet. One way to do this is to send IP directed broadcast packets targeted at the hosts in a particular target subnet.

The network forwards IP directed broadcast packets as if they were unicast packets. When the IP directed broadcast packet is received by a VLAN that is enabled for **targeted-broadcast**, the switch broadcasts the packet to all the hosts in its subnet.

In this topology (see [Figure 5 on page 55](#)), a host is connected to an interface on a switch to manage the clients in subnet **10.1.2.1/24**. When the switch receives a packet with the broadcast IP address of the target subnet as its destination address, it forwards the packet to the subnet's Layer 3 interface and broadcasts it to all the hosts within the subnet.

Figure 5: Topology for IP Directed Broadcast

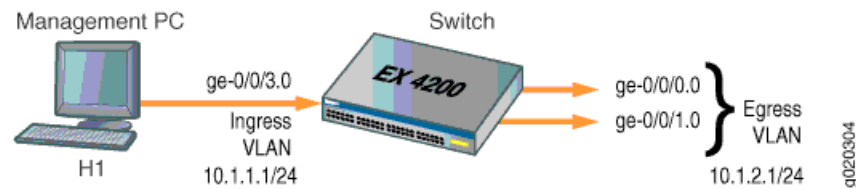


Table 11 on page 55 shows the settings of the components in this example.

Table 11: Components of the IP Directed Broadcast Topology

Property	Settings
Ingress VLAN name	v0
Ingress VLAN IP address	10.1.1.1/24
Egress VLAN name	v1
Egress VLAN IP address	10.1.2.1/24
Interfaces in VLAN v0	ge-0/0/3.0
Interfaces in VLAN v1	ge-0/0/0.0 and ge-0/0/1.0

Configuration

To configure IP directed broadcast on a subnet to enable remote management of its hosts:

CLI Quick Configuration

To quickly configure the switch to accept IP directed broadcasts targeted at subnet 10.1.2.1/24, copy the following commands and paste them into the switch's terminal window:

```
[edit]
set interfaces ge-0/0/0.0 family ethernet-switching vlan members v1
set interfaces ge-0/0/1.0 family ethernet-switching vlan members v1
set interfaces vlan.1 family inet address 10.1.2.1/24
set interfaces ge-0/0/3.0 family ethernet-switching vlan members v0
set interfaces vlan.0 family inet address 10.1.1.1/24
set vlans v1 l3-interface vlan.1
set vlans v0 l3-interface vlan.0
set interfaces vlan.1 family inet targeted-broadcast
```

Step-by-Step Procedure

To configure the switch to accept IP directed broadcasts targeted at subnet 10.1.2.1/24:

1. Add logical interface **ge-0/0/0.0** to VLAN **v1**:

```
[edit interfaces]
user@switch# set ge-0/0/0.0 family ethernet-switching vlan members v1
```
2. Add logical interface **ge-0/0/1.0** to VLAN **v1**:

```
[edit interfaces]
user@switch# set ge-0/0/1.0 family ethernet-switching vlan members v1
```
3. Configure the IP address for the egress VLAN, **v1**:

- ```
[edit interfaces]
user@switch# set vlan.1 family inet address 10.1.2.1/24
```
4. Add logical interface **ge-0/0/3.0** to VLAN **v0**:
 

```
[edit interfaces]
user@switch# set ge-0/0/3.0 family ethernet-switching vlan members v0
```
  5. Configure the IP address for the ingress VLAN:
 

```
[edit interfaces]
user@switch# set vlan.0 family inet address 10.1.1.1/24
```
  6. To route traffic between the ingress and egress VLANs, associate a Layer 3 interface with each VLAN:
 

```
[edit vlans]
user@switch# set v1 l3-interface (VLANs)vlan.1
user@switch# set v0 l3-interface vlan.0
```
  7. Enable the Layer 3 interface for the egress VLAN to receive IP directed broadcasts:
 

```
[edit interfaces]
user@switch# set vlan.1 family inet targeted-broadcast
```

**Results** Check the results:

```
user@switch# show
interfaces {
 ge-0/0/0 {
 unit 0 {
 family ethernet-switching {
 vlan {
 members v1;
 }
 }
 }
 }
 ge-0/0/1 {
 unit 0 {
 family ethernet-switching {
 vlan {
 members v1;
 }
 }
 }
 }
 ge-0/0/3 {
 unit 0 {
 family ethernet-switching {
 vlan {
 members v0;
 }
 }
 }
 }
 vlan {
 unit 0 {
 family inet {
 targeted-broadcast;
 address 10.1.1.1/24;
 }
 }
 }
}
```

```

 }
 unit 1 {
 family inet {
 targeted-broadcast;
 address 10.1.2.1/24;
 }
 }
}
vlands {
 default;
 v0 {
 l3-interface vland.0;
 }
 v1 {
 l3-interface vland.1;
 }
}
}

```

**Related Documentation**

- [Configuring IP Directed Broadcast \(CLI Procedure\) on page 126](#)

## Example: Configuring Multicast Load Balancing for Use with Aggregated 10-Gigabit Ethernet Interfaces on EX8200 Switches

EX8200 switches support multicast load balancing on link aggregation groups (LAGs). Multicast load balancing evenly distributes Layer 3 routed multicast traffic over the LAGs. You can aggregate up to twelve 10-gigabit Ethernet links to form a 120-gigabit virtual link or LAG. The MAC client can treat this virtual link as if it were a single link to increase bandwidth, provide graceful degradation as link failures occur, and increase availability. On EX8200 switches, multicast load balancing is enabled by default. However, if it is explicitly disabled, you can reenale it. .



**NOTE:** An interface with an already configured IP address cannot form part of the LAG.



**NOTE:** Only EX8200 standalone switches with 10-gigabit links support multicast load balancing. Virtual Chassis does not support multicast load balancing.

This example shows how to configure a LAG and reenale multicast load balancing:

- [Requirements on page 58](#)
- [Overview and Topology on page 58](#)
- [Configuration on page 59](#)
- [Verification on page 61](#)

## Requirements

This example uses the following hardware and software components:

- Two EX8200 switches, one used as the access switch and one used as the distribution switch
- Junos OS Release 12.2 or later for EX Series switches

Before you begin:

- Configure four 10-gigabit interfaces on the EX8200 distribution switch: xe-0/1/0, xe-1/1/0, xe-2/1/0, and xe-3/1/0. See [“Configuring Gigabit Ethernet Interfaces \(CLI Procedure\)” on page 64](#).

## Overview and Topology

Multicast load balancing uses one of seven hashing algorithms to balance traffic between the individual 10-gigabit links in the LAG. For a description of the hashing algorithms, see [multicast-loadbalance](#). The default hashing algorithm is `crc-sgip`. You can experiment with the different hashing algorithms until you determine the one that best balances your Layer 3 routed multicast traffic.

When a link larger than 10 gigabits is needed on an EX8200 switch, you can combine up to twelve 10-gigabit links to create more bandwidth. This example uses the link aggregation feature to combine four 10-gigabit links into a 40-gigabit link on the distribution switch. In addition, multicast load balancing is enabled to ensure even distribution of Layer 3 routed multicast traffic on the 40-gigabit link. In the sample topology illustrated in [Figure 6 on page 59](#), an EX8200 switch in the distribution layer is connected to an EX8200 switch in the access layer.



**NOTE:** Link speed is automatically determined based on the size of the LAG configured. For example, if a LAG is composed of four 10-gigabit links, the link speed is 40 gigabits per second).

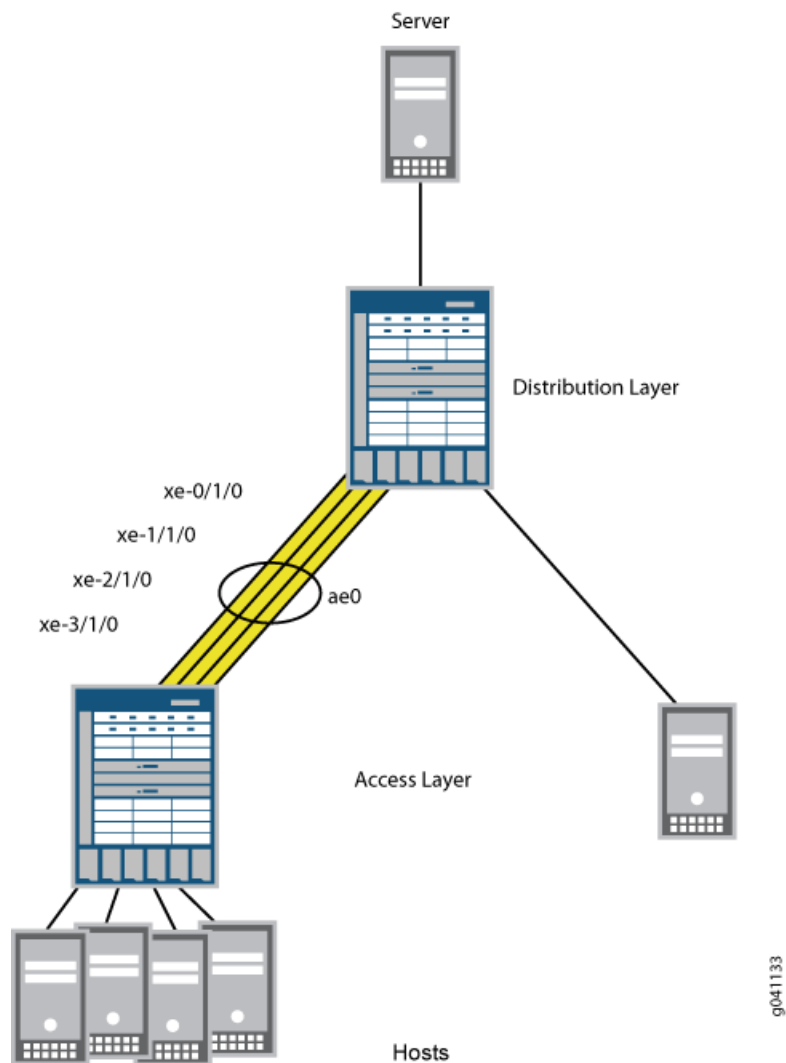
---



**NOTE:** The default hashing algorithm, `crc-sgip`, involves a cyclic redundancy check of both the multicast packet source and group IP addresses.

---

Figure 6: 40-Gigabit LAG Composed of Four 10-Gigabit Links



You will configure a LAG on each switch and reenables multicast load balancing. When reenabled, multicast load balancing will automatically take effect on the LAG, and the speed is set to 10 gigabits per second for each link in the LAG. Link speed for the 40-gigabit LAG is automatically set to 40 gigabits per second.

## Configuration


### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set chassis aggregated-devices ethernet device-count 1
set interfaces ae0 aggregated-ether-options minimum-links 1
set interfaces xe-0/1/0 ether-options 802.3ad ae0
set interfaces xe-1/1/0 ether-options 802.3ad ae0
set interfaces xe-2/1/0 ether-options 802.3ad ae0
```

- Step-by-Step Procedure**
- To configure a LAG and reenable multicast load balancing:
- Specify the number of aggregated Ethernet interfaces to be created:
 


```
[edit chassis]
user@switch# set aggregated-devices ethernet device-count 1
```
  - Specify the minimum number of links for the aggregated Ethernet interface (aex), that is, the LAG, to be labeled up:
 

 **NOTE:** By default, only one link needs to be up for the LAG to be labeled up.

```
[edit interfaces]
user@switch# set ae0 aggregated-ether-options minimum-links 1
```
  - Specify the four members to be included within the LAG:
 

```
[edit interfaces]
user@switch# set xe-0/1/0 ether-options 802.3ad ae0
user@switch# set xe-1/1/0 ether-options 802.3ad ae0
user@switch# set xe-2/1/0 ether-options 802.3ad ae0
user@switch# set xe-3/1/0 ether-options 802.3ad ae0
```
  - Reenable multicast load balancing:
 

```
[edit chassis]
user@switch# set multicast-loadbalance
```

 **NOTE:** You do not need to set link speed the way you do for LAGs that do not use multicast load balancing. Link speed is automatically set to 40 gigabits per second on a 40-gigabit LAG.
  - You can optionally change the value of the **hash-mode** option in the **multicast-loadbalance** statement to try different algorithms until you find the one that best distributes your Layer 3 routed multicast traffic.
 

If you change the hashing algorithm when multicast load balancing is disabled, the new algorithm takes effect after you reenable multicast load balancing.

**Results** Check the results of the configuration:

```
user@switch> show configuration
chassis
aggregated-devices {
 ethernet {
 device-count 1;
 }
}
multicast-loadbalance {
 hash-mode crc-gip;
```

```

}

interfaces
xe-0/1/0 {
ether-options {
802.3ad ae0;
}
}
xe-1/1/0 {
ether-options {
802.3ad ae0;
}
}
xe-2/1/0 {
ether-options {
802.3ad ae0;
}
}
xe-3/1/0 {
ether-options {
802.3ad ae0;
}
}
ae0 {
aggregated-ether-options {
minimum-links 1;
}
}
}

```

## Verification

To confirm that the configuration is working properly, perform these tasks:

- [Verifying the Status of a LAG Interface on page 61](#)
- [Verifying Multicast Load Balancing on page 62](#)

### Verifying the Status of a LAG Interface

**Purpose** Verify that a link aggregation group (LAG) (**ae0**) has been created on the switch.

**Action** Verify that the **ae0** LAG has been created:

```
user@switch> show interfaces ae0 terse
```

| Interface | Admin | Link | Proto | Local         | Remote |
|-----------|-------|------|-------|---------------|--------|
| ae0       | up    | up   |       |               |        |
| ae0.0     | up    | up   | inet  | 10.10.10.2/24 |        |

**Meaning** The interface name *aex* indicates that this is a LAG. *A* stands for aggregated, and *E* stands for Ethernet. The number differentiates the various LAGs.

### Verifying Multicast Load Balancing

---

**Purpose** Check that traffic is load-balanced equally across paths.

**Action** Verify load balancing across the four interfaces:

```
user@switch> monitor interface traffic

Bytes=b, Clear=c, Delta=d, Packets=p, Quit=q or ESC, Rate=r, Up=^U, Down=^D
ibmoem02-re1 Seconds: 3 Time: 16:06:14

Interface Link Input packets (pps) Output packets (pps)
xe-0/1/0 Up 2058834 (10) 7345862 (19)
xe-1/1/0 Up 2509289 (9) 6740592 (21)
xe-2/1/0 Up 8625688 (90) 10558315 (20)
xe-3/1/0 Up 2374154 (23) 71494375 (9)
```

**Meaning** The interfaces should be carrying approximately the same amount of traffic.

- Related Documentation**
- [Configuring Multicast Load Balancing for Use with Aggregated 10-Gigabit Ethernet Links on EX8200 Switches \(CLI Procedure\) on page 132](#)
  - [Understanding Multicast Load Balancing on Aggregated 10-Gigabit Links for Routed Multicast Traffic on EX8200 Switches on page 24](#)

## CHAPTER 3

# Configuration Tasks

- [Configuring Gigabit Ethernet Interfaces \(CLI Procedure\) on page 64](#)
- [Configuring Gigabit Ethernet Interfaces \(J-Web Procedure\) on page 68](#)
- [Port Role Configuration with the J-Web Interface \(with CLI References\) on page 75](#)
- [Adding a Logical Unit Description to the Configuration on page 79](#)
- [Disabling a Physical Interface on page 80](#)
- [Disabling a Logical Interface on page 82](#)
- [Configuring Flow Control on page 82](#)
- [Configuring the Interface Address on page 84](#)
- [Configuring the Interface Bandwidth on page 86](#)
- [Configuring the Media MTU on page 87](#)
- [Setting the Protocol MTU on page 88](#)
- [Interface Ranges on page 89](#)
- [Configuring Accounting for the Physical Interface on page 98](#)
- [Configuring Accounting for the Logical Interface on page 101](#)
- [Configuring Ethernet Loopback Capability on page 103](#)
- [Configuring Gratuitous ARP on page 104](#)
- [Configuring Static ARP Table Entries For Mapping IP Addresses to MAC Addresses on page 105](#)
- [Disabling the Transmission of Redirect Messages on an Interface on page 106](#)
- [Configuring Restricted and Unrestricted Proxy ARP on page 107](#)
- [Enabling or Disabling SNMP Notifications on Logical Interfaces on page 108](#)
- [Enabling or Disabling SNMP Notifications on Physical Interfaces on page 108](#)
- [Configuring Aggregated Ethernet Links \(CLI Procedure\) on page 109](#)
- [Configuring Aggregated Ethernet Interfaces \(J-Web Procedure\) on page 110](#)
- [Configuring Aggregated Ethernet LACP \(CLI Procedure\) on page 113](#)
- [Configuring LACP Link Protection of Aggregated Ethernet Interfaces \(CLI Procedure\) on page 114](#)
- [Configuring Aggregated Ethernet Link Protection on page 118](#)

- [Configuring Aggregated Ethernet Link Speed on page 120](#)
- [Configuring Aggregated Ethernet Minimum Links on page 121](#)
- [Configuring Tagged Aggregated Ethernet Interfaces on page 122](#)
- [Configuring a Layer 3 Subinterface \(CLI Procedure\) on page 123](#)
- [Configuring Unicast RPF \(CLI Procedure\) on page 123](#)
- [Disabling Unicast RPF \(CLI Procedure\) on page 125](#)
- [Configuring IP Directed Broadcast \(CLI Procedure\) on page 126](#)
- [Tracing Operations of an Individual Router or Switch Interface on page 127](#)
- [Tracing Operations of the Interface Process on page 128](#)
- [Setting the Mode on an SFP+ or SFP+ MACSec Uplink Module \(CLI Procedure\) on page 129](#)
- [Configuring the Media Type on Dual-Purpose Uplink Ports \(CLI Procedure\) on page 130](#)
- [Configuring Generic Routing Encapsulation Tunneling \(CLI Procedure\) on page 130](#)
- [Configuring Multicast Load Balancing for Use with Aggregated 10-Gigabit Ethernet Links on EX8200 Switches \(CLI Procedure\) on page 132](#)
- [Configuring Energy Efficient Ethernet on Interfaces \(CLI Procedure\) on page 133](#)
- [Damping Shorter Physical Interface Transitions on page 135](#)
- [Configuring Reflective Relay on page 136](#)

## Configuring Gigabit Ethernet Interfaces (CLI Procedure)

---



**NOTE:** This task uses Junos OS for EX Series switches that does not support the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that supports ELS, see *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*. For ELS details, see *Getting Started with Enhanced Layer 2 Software*.

An Ethernet interface must be configured for optimal performance in a high-traffic network. EX Series switches include a factory default configuration that:

- Enables all the network interfaces on the switch
- Sets a default port mode (access)
- Sets default link settings

- Specifies a logical unit (**unit 0**) and assigns it to **family ethernet-switching** (except on EX8200 switches and Virtual Chassis)
- Specifies Rapid Spanning Tree Protocol (RSTP) and Link Layer Discovery Protocol (LLDP)

This topic describes:

- [Configuring VLAN Options and Port Mode on page 65](#)
- [Configuring the Link Settings on page 66](#)
- [Configuring the IP Options on page 68](#)

## Configuring VLAN Options and Port Mode

By default, when you boot a switch and use the factory default configuration, or when you boot the switch and do not explicitly configure a port mode, all interfaces on the switch are in access mode and accept only untagged packets from the VLAN named **default**. You can optionally configure another VLAN and use that instead of **default**. You can also configure a port to accept untagged packets from the user-configured VLAN. For details on this concept (native VLAN), see *Understanding Bridging and VLANs on EX Series Switches*.

If you are connecting either a desktop phone, wireless access point or a security camera to a Power over Ethernet (PoE) port, you can configure some parameters for the PoE interface. PoE interfaces are enabled by default. For detailed information on PoE settings, see *Configuring PoE on EX Series Switches (CLI Procedure)*.

If you are connecting a device to other switches and to routers on the LAN, you need to assign the interface to a logical port and configure the logical port as a trunk port. See [“Port Role Configuration with the J-Web Interface \(with CLI References\)” on page 75](#) for more information about port configuration.

If you are connecting to a server that contains virtual machines and a VEPA for packet aggregation from those virtual machines, configure the port as a tagged-access port. See *Understanding Bridging and VLANs on EX Series Switches* for more information about tagged access.

To configure a Gigabit Ethernet interface or 10-Gigabit Ethernet interface for trunk port mode:

```
[edit]
user@switch# set interfaces interface-name unit logical-unit-number family ethernet-switching
port-mode trunk
```

To configure a Gigabit Ethernet interface or 10-Gigabit Ethernet interface for tagged-access port mode:

```
[edit]
user@switch# set interfaces interface-name unit logical-unit-number family ethernet-switching
port-mode tagged-access
```

## Configuring the Link Settings

EX Series switches include a factory default configuration that enables interfaces with the following link settings:

- All Gigabit Ethernet interfaces are set to **auto-negotiation**.
- The speed for Gigabit Ethernet interfaces is set to **auto**, allowing the interface to operate at 10m, 100m, or 1g. The link operates at the highest possible speed, depending on the capabilities of the remote end.
- The flow control for Gigabit Ethernet interfaces and 10-Gigabit Ethernet interfaces is set to **enabled**.
- The link mode is set to **auto**, allowing the interface to operate as either full duplex or half duplex. The link operates as full duplex unless this mode is not supported at the remote end.
- The 10-Gigabit Ethernet fiber interfaces default to **no auto-negotiation**. The default speed is 10g and the default link mode is full duplex.

To configure the link settings:

- Set link settings for a Gigabit Ethernet interface:

```
[edit]
user@switch# set interfaces ge-fpc/pic/port ether-options
```

- Set link settings for a 10-Gigabit Ethernet interface:

```
[edit]
user@switch# set interfaces xe-fpc/pic/port ether-options
```



**NOTE:** On EX Series switches, *fpc* can have the following values:

- On an EX2200 switch, an EX3200 switch, a standalone EX3300 switch, a standalone EX4200 switch, and a standalone EX4500 switch, FPC refers to the switch itself. The FPC number is always 0 on these switches.
- On an EX3300 Virtual Chassis, an EX4200 Virtual Chassis, an EX4500 Virtual Chassis, or a mixed EX4200 and EX4500 Virtual Chassis, the FPC number indicates the member ID of the switch within the Virtual Chassis.
- On an EX6200 switch and a standalone EX8200 switch, the FPC number indicates the slot number of the line card that contains the physical interface. On an EX6200 switch, the FPC number also indicates the slot number of the Switch Fabric and Routing Engine (SRE) module that contains the uplink port.
- On an EX8200 Virtual Chassis, the FPC number indicates the slot number of the line card on the Virtual Chassis. The line card slots on Virtual Chassis member 0 are numbered 0 through 15; on member 1, they are numbered 16 through 31, and so on.

*pic* can have the following values:

- On EX2200, EX3200, EX3300, EX4200, and EX4500 switches, the PIC number is 0 for all built-in interfaces (interfaces that are not an uplink port).
- On EX2200, EX3200, and EX4200 switches, the PIC number is 1 for uplink ports.
- On EX4500 switches, the PIC number is 1 for uplink ports on the left-hand uplink module and 2 for uplink ports on the right-hand uplink module.
- On EX6200 and EX8200 switches, the PIC number is always 0.

The **ether-options** statement allows you to modify the configuration:

- **802.3ad**—Specify an aggregated Ethernet bundle. See “[Configuring Aggregated Ethernet Links \(CLI Procedure\)](#)” on page 109.
- **auto-negotiation**—Enable or disable autonegotiation of flow control, link mode, and speed.
- **flow-control**—Enable or disable flow control.

- **link-mode**—Specify **full-duplex**, **half-duplex**, or **automatic**.
- **loopback**—Enable or disable loopback mode.
- **speed**—Specify **10m**, **100m**, **1g**, or **autonegotiation**.

## Configuring the IP Options

To specify an IP address for the logical unit using IPv4:

[edit]

```
user@switch# set interfaces interface-name unit logical-unit-number family inet address ip-address
```

To specify an IP address for the logical unit using IPv6:

[edit]

```
user@switch# set interfaces interface-name unit logical-unit-number family inet6 address ip-address
```



**NOTE:** Access interfaces on EX2200, EX3200, EX3300, EX4200, and EX4500 switches are set to family ethernet-switching by default. You might have to delete this or another user-configured family setting before changing the setting to family inet or family inet6.

### Related Documentation

- [Configuring Gigabit Ethernet Interfaces \(J-Web Procedure\) on page 68](#)
- [Monitoring Interface Status and Traffic on page 277](#)
- [show interfaces ge- on page 349](#)
- [show interfaces xe- on page 368](#)
- [Understanding Interface Naming Conventions on EX Series Switches on page 6](#)

---

## Configuring Gigabit Ethernet Interfaces (J-Web Procedure)

You can configure specific properties on your Ethernet interface to ensure optimal performance of your network in a high-traffic environment.

To configure properties on a Gigabit Ethernet interface, a 10-Gigabit Ethernet interface, and a 40-Gigabit Ethernet interface on an EX Series switch:

1. Select **Interfaces > Ports**.

The page that is displayed lists Gigabit Ethernet, 10-Gigabit Ethernet interfaces, and 40-Gigabit Ethernet interfaces, and their link statuses.



**NOTE:** After you make changes to the configuration on this page, you must commit the changes immediately for them to take effect. To commit all changes to the active configuration, select **Commit Options > Commit**. See *Using the Commit Options to Commit Configuration Changes (J-Web Procedure)* for details about all commit options.

2. Select the interface you want to configure. For an EX8200 Virtual Chassis configuration, select the member and the FPC slot if the interface you want to configure is not listed under **Ports** in the top table on the page.

Details for the selected interface, such as administrative status, link status, speed, duplex, and flow control, are displayed in the **Details of port** table on the page.



**NOTE:** You can select multiple interfaces and modify their settings at the same time. However, while doing this, you cannot modify the IP address or enable or disable the administrative status of the selected interfaces.



**NOTE:** In the J-Web interface, you cannot configure interface ranges and interface groups.

3. Click **Edit** and select the set of options you want to configure first:

- Port Role—Enables you to assign a profile for the selected interface.



**NOTE:** When you select a particular port role, preconfigured port security parameters are set for the VLAN that the interface belongs to. For example, if you select the port role **Desktop**, the port security options **examine-dhcp** and **arp-inspection** are enabled on the VLAN that the interface belongs to. If there are interfaces in the VLAN that have static IP addresses, those interfaces might lose connectivity because those static IP addresses might not be present in the DHCP pool. Therefore, when you select a port role, ensure that the corresponding port security settings for the VLAN are applicable to the interface.

For basic information about port security features such as DHCP snooping (CLI option **examine-dhcp**) or dynamic ARP inspection (DAI) (CLI option **arp-inspection**), see *Configuring Port Security (J-Web Procedure)*. For detailed descriptions of port security features, see the Port Security topics in the EX Series documentation at <http://www.juniper.net/techpubs/>.

Click **Details** to view the configuration parameters for the selected port role.

- VLAN—Enables you to configure VLAN options for the selected interface.
  - Link—Enables you to modify the following link options for the selected interface:
    - Speed
    - MTU
    - Autonegotiation
    - Flow Control
    - Duplex
    - Media Type
  - IP—Enables you to configure an IP address for the interface.
4. Configure the interface by configuring options in the selected option set. See [Table 12 on page 71](#) for details of the options.
  5. Repeat Steps 3 and 4 for the remaining option sets that you want to configure for the interface.



**NOTE:** To enable or disable the administrative status of a selected interface, click **Enable Port** or **Disable Port**.

Table 12: Port Edit Options

| Field             | Function                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | Your Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Port Role Options |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Port Role         | <p>Specifies a profile (role) to assign to the interface.</p> <p><b>NOTE:</b> After a port role is configured on the interface, you cannot specify VLAN options or IP options.</p> <p><b>NOTE:</b> Port roles are not supported by the <code>et</code> interfaces (40-Gigabit Ethernet interfaces) on EX4300 and EX4550 switches.</p> <p><b>NOTE:</b> Only the following port roles can be applied on EX8200 switch interfaces:</p> <ul style="list-style-type: none"> <li>• Default</li> <li>• Layer 2 uplink</li> <li>• Routed uplink</li> </ul>                                                                                                                             |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Default           | <p>Applies the default role.</p> <p>The interface family is set to <b>ethernet-switching</b>, port mode is set to <b>access</b>, and RSTP is enabled.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | <ol style="list-style-type: none"> <li>1. Click <b>Details</b> to view CLI commands for this role.</li> <li>2. Click <b>OK</b>.</li> </ol>                                                                                                                                                                                                                                                                                                                                                                                 |
| Desktop           | <p>Applies the desktop role.</p> <p>The interface family is set to <b>ethernet-switching</b>, port mode is set to <b>access</b>, RSTP is enabled with the <b>edge</b> and <b>point-to-point</b> options, and port security parameters (MAC limit =1; dynamic ARP inspection and DHCP snooping enabled) are set.</p> <p>To enable security configuration, select the <b>Enable Security Configuration</b> check box. The <b>forwarding-options dhcp-security-arp-inspection</b> will be configured.</p>                                                                                                                                                                         | <ol style="list-style-type: none"> <li>1. Select an existing VLAN configuration or type the name of a new VLAN configuration to be associated with the interface.</li> <li>2. Click <b>Details</b> to view CLI commands for this role.</li> <li>3. Click <b>OK</b>.</li> </ol>                                                                                                                                                                                                                                             |
| Desktop and Phone | <p>Applies the desktop and phone role.</p> <p>The interface family is set to <b>ethernet-switching</b>, port mode is set to <b>access</b>, port security parameters (MAC limit =1; dynamic ARP Inspection and DHCP snooping enabled) are set, and recommended class-of-service (CoS) parameters are specified for forwarding classes, schedulers, and classifiers. See <a href="#">Table 13 on page 74</a> for more CoS information.</p> <p>To enable security configuration, select the <b>Enable Security Configuration</b> check box. The <b>forwarding-options dhcp-security</b> groups and <b>forwarding-options dhcp-security-arp-inspection</b> will be configured.</p> | <ol style="list-style-type: none"> <li>1. Select an existing VLAN configuration or type the name of a new VLAN configuration to be associated with the interface.</li> </ol> <p>You can also select an existing VoIP VLAN configuration or a new VoIP VLAN configuration to be associated with the interface.</p> <p><b>NOTE:</b> VoIP is not supported on EX8200 switches.</p> <ol style="list-style-type: none"> <li>2. Click <b>Details</b> to view CLI commands for this role.</li> <li>3. Click <b>OK</b>.</li> </ol> |

Table 12: Port Edit Options (*continued*)

| Field                                                                                                      | Function                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | Your Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Wireless Access Point                                                                                      | <p>Applies the wireless access point role.</p> <p>The interface family is set to <b>ethernet-switching</b>, port mode is set to <b>access</b>, and RSTP is enabled with the <b>edge</b> and <b>point-to-point</b> options.</p>                                                                                                                                                                                                                                                            | <ol style="list-style-type: none"> <li>1. Select an existing VLAN configuration or type the name of a new VLAN configuration to be associated with the interface. Type the VLAN ID for a new VLAN.</li> <li>2. Click <b>Details</b> to view CLI commands for this role.</li> <li>3. Click <b>OK</b>.</li> </ol>                                                                                                                                                                                                                                                                                                                                                                                                 |
| Routed Uplink                                                                                              | <p>Applies the routed uplink role.</p> <p>The interface family is set to <b>inet</b>, and recommended CoS parameters are set for schedulers and classifiers. See <a href="#">Table 13 on page 74</a> for more CoS information.</p>                                                                                                                                                                                                                                                        | <p>To specify an IPv4 address:</p> <ol style="list-style-type: none"> <li>1. Select the <b>IPv4 address</b> check box.</li> <li>2. Type an IP address—for example: <b>10.10.10.10</b>.</li> <li>3. Enter the subnet mask or address prefix. For example, 24 bits represents <b>255.255.255.0</b>.</li> <li>4. Click <b>OK</b>.</li> </ol> <p>To specify an IPv6 address:</p> <ol style="list-style-type: none"> <li>1. Select the <b>IPv6 address</b> check box.</li> <li>2. Type an IP address—for example: <b>2001:ab8:85a3::8a2e:370:7334</b>.</li> <li>3. Enter the subnet mask or address prefix.</li> <li>4. Click <b>OK</b>.</li> </ol> <p><b>NOTE:</b> IPv6 is not supported on EX2200 VC switches.</p> |
| Layer 2 Uplink                                                                                             | <p>Applies the Layer 2 uplink role.</p> <p>The interface family is set to <b>ethernet-switching</b>, port mode is set to <b>trunk</b>, RSTP is enabled with the <b>point-to-point</b> option, and trusted DHCP is configured for port security.</p> <p>To enable security configuration, select the <b>Enable Security Configuration</b> check box. The <b>forwarding-options dhcp-security</b> groups and <b>forwarding-options dhcp-security-arp-inspection</b> will be configured.</p> | <ol style="list-style-type: none"> <li>1. For this port role, you can select a VLAN member and associate a native VLAN with the interface.</li> <li>2. Click <b>Details</b> to view CLI commands for this role.</li> <li>3. Click <b>OK</b>.</li> </ol>                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| None                                                                                                       | Specifies that no port role is configured for the selected interface.                                                                                                                                                                                                                                                                                                                                                                                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>NOTE:</b> For an EX8200 switch, dynamic ARP inspection and DHCP snooping parameters are not configured. |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| VLAN Options                                                                                               |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

Table 12: Port Edit Options (*continued*)

| Field                   | Function                                                                                                                                                                                           | Your Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|-------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Port Mode               | Specifies the mode of operation for the interface: trunk or access.                                                                                                                                | <p>If you select <b>Trunk</b>, you can:</p> <ol style="list-style-type: none"> <li>1. Click <b>Add</b> to add a VLAN member.</li> <li>2. Select the VLAN and click <b>OK</b>.</li> <li>3. (Optional) Associate a native VLAN with the interface.</li> <li>4. Click <b>OK</b>.</li> </ol> <p>If you select <b>Access</b>, you can:</p> <ol style="list-style-type: none"> <li>1. Select the VLAN member to be associated with the interface.</li> <li>2. (Optional) Associate a VoIP VLAN with the interface. Only a VLAN with a VLAN ID can be associated as a VoIP VLAN.</li> </ol> <p><b>NOTE:</b> VoIP is not supported on EX8200 switches.</p> <ol style="list-style-type: none"> <li>3. Click <b>OK</b>.</li> </ol> |
| Link Options            |                                                                                                                                                                                                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| MTU (bytes)             | Specifies the maximum transmission unit size (MTU) for the interface.                                                                                                                              | Type a value from <b>256</b> through <b>9216</b> . The default MTU size for Gigabit Ethernet interfaces is <b>1514</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| Speed                   | Specifies the speed for the mode.                                                                                                                                                                  | <p>Select one of the following values: <b>10 Mbps</b>, <b>100 Mbps</b>, <b>1000 Mbps</b>, or <b>Auto-Negotiation</b>.</p> <p><b>NOTE:</b> EX4300 switches supports <b>Auto-Negotiation 10M-100M</b> apart from the values mentioned above.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Duplex                  | Specifies the link mode.                                                                                                                                                                           | <p>Select one: <b>automatic</b>, <b>half</b>, or <b>full</b>.</p> <p><b>NOTE:</b> Link mode <b>half</b> is not supported on EX4300 switches.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Description             | <p>Describes the link.</p> <p><b>NOTE:</b> If the interface is part of a link aggregation group (LAG), only the <b>Description</b> option is enabled. Other Port Edit options are unavailable.</p> | Enter a brief description for the link.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Enable Auto Negotiation | Enables or disables autonegotiation.                                                                                                                                                               | Select the check box to enable autonegotiation, or clear the check box to disable it. By default, autonegotiation is enabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Enable Flow Control     | Enables or disables flow control.                                                                                                                                                                  | Select the check box to enable flow control to regulate the amount of traffic sent out of the interface, or clear the check box to disable flow control and permit unrestricted traffic. Flow control is enabled by default.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

Table 12: Port Edit Options (*continued*)

| Field        | Function                                                                                                                                                          | Your Action                                                                                                                                                                                                                                                                                                                                                                   |
|--------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Media Type   | Specifies the media type selected.                                                                                                                                | Select the check box to enable the media type. Then select <b>Copper</b> or <b>Fiber</b> .                                                                                                                                                                                                                                                                                    |
| IP Options   |                                                                                                                                                                   |                                                                                                                                                                                                                                                                                                                                                                               |
| IPv4 Address | Specifies an IPv4 address for the interface.<br><br><b>NOTE:</b> If the IPv4 Address check box is cleared, the interface still belongs to the <b>inet</b> family. | <ol style="list-style-type: none"> <li>1. Select the <b>IPv4 address</b> check box to specify an IPv4 address.</li> <li>2. Type an IP address—for example: <b>10.10.10.10</b>.</li> <li>3. Enter the subnet mask or address prefix. For example, 24 bits represents <b>255.255.255.0</b>.</li> <li>4. Click <b>OK</b>.</li> </ol>                                             |
| IPv6 Address | Specifies an IPv6 address for the interface.<br><br><b>NOTE:</b> If the IPv6 Address check box is cleared, the interface still belongs to the <b>inet</b> family. | <ol style="list-style-type: none"> <li>1. Select the <b>IPv6 address</b> check box to specify an IPv6 address.</li> <li>2. Type an IP address—for example: <b>2001:ab8:85a3::8a2e:370:7334</b>.</li> <li>3. Enter the subnet mask or address prefix.</li> <li>4. Click <b>OK</b>.</li> </ol> <p><b>NOTE:</b> IPv6 address is not supported on EX2200 and EX4500 switches.</p> |

Table 13: Recommended CoS Settings for Port Roles

| CoS Parameter         | Recommended Settings                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|-----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Forwarding Classes    | <p>There are four forwarding classes:</p> <ul style="list-style-type: none"> <li>• <b>voice</b>—Queue number is set to 7.</li> <li>• <b>expedited-forwarding</b>—Queue number is set to 5.</li> <li>• <b>assured-forwarding</b>—Queue number is set to 1.</li> <li>• <b>best-effort</b>—Queue number is set to 0.</li> </ul>                                                                                                                                                                                                                                                     |
| Schedulers            | <p>The schedulers and their settings are:</p> <ul style="list-style-type: none"> <li>• Strict-priority—Transmission rate is set to 10 percent and buffer size to 5 percent.</li> <li>• Expedited-scheduler—Transmission rate is set to 30 percent, buffer size to 30 percent, and priority to <b>low</b>.</li> <li>• Assured-scheduler—Transmission rate is set to 25 percent, buffer size to 25 percent, and priority to <b>low</b>.</li> <li>• Best-effort scheduler—Transmission rate is set to 35 percent, buffer size to 40 percent, and priority to <b>low</b>.</li> </ul> |
| Scheduler maps        | When a desktop and phone, routed uplink, or Layer 2 uplink role is applied on an interface, the forwarding classes and schedulers are mapped using the scheduler map.                                                                                                                                                                                                                                                                                                                                                                                                            |
| ieee-802.1 classifier | Imports the default <b>ieee-802.1</b> classifier configuration and sets the loss priority to <b>low</b> for the code point 101 for the <b>voice</b> forwarding class.                                                                                                                                                                                                                                                                                                                                                                                                            |

Table 13: Recommended CoS Settings for Port Roles (*continued*)

| CoS Parameter                | Recommended Settings                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| dscp classifier              | Imports the default <b>dscp</b> classifier configuration and sets the loss priority to <b>low</b> for the code point 101110 for the <b>voice</b> forwarding class.                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b> | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 64</a></li> <li>• <a href="#">Configuring Gigabit Ethernet Interfaces (CLI Procedure)</a></li> <li>• <a href="#">Monitoring Interface Status and Traffic on page 277</a></li> <li>• <a href="#">EX Series Switches Interfaces Overview on page 3</a></li> <li>• <a href="#">Junos OS CoS for EX Series Switches Overview</a></li> <li>• <a href="#">Understanding Interface Naming Conventions on EX Series Switches on page 6</a></li> </ul> |

## Port Role Configuration with the J-Web Interface (with CLI References)

When you configure Gigabit Ethernet interface properties with the J-Web interface (Configure > Interfaces) you can optionally select pre-configured port roles for those interfaces. When you select a role from the **Port Role** field and apply it to a port, the J-Web interface modifies the switch configuration using CLI commands. [Table 14 on page 75](#) lists the CLI commands applied for each port role.



**NOTE:** If there is an existing port role configuration, it is cleared before the new port role configuration is applied.

Table 14: Port Role Configuration Summary

| Configuration Description                                 | CLI Commands                                                                          |
|-----------------------------------------------------------|---------------------------------------------------------------------------------------|
| Default Port Role                                         |                                                                                       |
| Set the port role to <b>Default</b> .                     | <code>set interfaces <i>interface</i> apply-macro juniper-port-profile Default</code> |
| Set port family to <b>ethernet-switching</b> .            | <code>set interfaces <i>interface</i> unit 0 family ethernet-switching</code>         |
| Set port mode to <b>access</b> .                          | <code>port-mode access</code>                                                         |
| Enable RSTP if redundant trunk groups are not configured. | <code>delete protocols rstp interface <i>interface</i> disable</code>                 |
| Disable RSTP if redundant trunk groups are configured.    | <code>set protocols rstp interface <i>interface</i> disable</code>                    |
| Desktop Port Role                                         |                                                                                       |

Table 14: Port Role Configuration Summary (*continued*)

| Configuration Description                                                                                                                                                            | CLI Commands                                                                                                                                                                                                                                                                                                          |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Set the port role to desktop.                                                                                                                                                        | <code>set interfaces <i>interface</i> apply-macro juniper-port-profile Desktop</code>                                                                                                                                                                                                                                 |
| Set VLAN if new VLAN is specified.                                                                                                                                                   | <code>set vlans &lt;<i>vlan name</i>&gt; vlan-id &lt;<i>vlan-id</i>&gt;</code>                                                                                                                                                                                                                                        |
| Set port family to <b>ethernet-switching</b> .<br>Set Port Mode to <b>Access</b> .                                                                                                   | <code>set interfaces <i>interface</i> unit 0 family ethernet-switching port-mode access</code>                                                                                                                                                                                                                        |
| Set VLAN if new VLAN is specified.                                                                                                                                                   | <code>set interfaces <i>interface</i> unit 0 family ethernet-switching vlan members <i>vlan-members</i></code>                                                                                                                                                                                                        |
| Set port security parameters.                                                                                                                                                        | <code>set ethernet-switching-options secure-access-port vlan MacTest arp-inspection</code>                                                                                                                                                                                                                            |
| Set RSTP protocol with <b>edge</b> option.                                                                                                                                           | <code>set protocols rstp interface <i>interface</i> edge</code>                                                                                                                                                                                                                                                       |
| RSTP protocol is disabled if redundant trunk groups are configured.                                                                                                                  | <code>set protocols rstp interface <i>interface</i> disable</code>                                                                                                                                                                                                                                                    |
| Desktop and Phone Port Role                                                                                                                                                          |                                                                                                                                                                                                                                                                                                                       |
| Set the port role to desktop and phone.                                                                                                                                              | <code>set interfaces <i>interface</i> apply-macro juniper-port-profile Desktop and Phone</code>                                                                                                                                                                                                                       |
| Set data VLAN if new VLAN is specified.<br>Set voice VLAN if new voice VLAN is specified.                                                                                            | <code>set vlans <i>vlan-name</i> vlan-id <i>vlan id</i></code>                                                                                                                                                                                                                                                        |
| Set port family to <b>ethernet-switching</b> .<br>Set Port Mode to <b>access</b> .                                                                                                   | <code>set interfaces <i>interface</i> unit 0 family ethernet-switching port-mode access</code>                                                                                                                                                                                                                        |
| Set data VLAN on port stanza.                                                                                                                                                        | <code>set interfaces <i>interface</i> unit 0 family ethernet-switching vlan members <i>vlan-members</i></code>                                                                                                                                                                                                        |
| Set port security parameters.                                                                                                                                                        | <code>set ethernet-switching-options secure-access-port vlan MacTest arp-inspection</code>                                                                                                                                                                                                                            |
| Set VOIP VLAN.                                                                                                                                                                       | <code>set ethernet-switching-options voip interface <i>interface</i>.0 vlan <i>vlan</i> <i>vlan name</i></code>                                                                                                                                                                                                       |
| Set class of service parameters<br><b>SCHEDULER_MAP=juniper-port-profile-map</b><br><b>IEEE_CLASSIFIER=juniper-ieee-classifier</b><br><b>DSCP_CLASSIFIER=juniper-dscp-classifier</b> | <code>set class-of-service interfaces <i>interfaces</i> scheduler-map juniper-port-profile-map<br/>set class-of-service interfaces <i>interface</i> unit 0 classifiers ieee-802.1 juniper_ieee_classifier<br/>set class-of-service interfaces <i>interface</i> unit 0 classifiers dscp juniper-dscp-classifier</code> |
| Set CoS Configuration                                                                                                                                                                | Refer <a href="#">Table 15 on page 78</a> for details.                                                                                                                                                                                                                                                                |
| Wireless Access Point Port Role                                                                                                                                                      |                                                                                                                                                                                                                                                                                                                       |
| Set the port role to wireless access point.                                                                                                                                          | <code>set interfaces <i>interface</i> apply-macro juniper-port-profile Wireless Access Point</code>                                                                                                                                                                                                                   |

Table 14: Port Role Configuration Summary (*continued*)

| Configuration Description                                                                                                                                                                                                      | CLI Commands                                                                                                                                                                                                                                                                                                                      |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Set VLAN on VLANs stanza.                                                                                                                                                                                                      | <code>set vlans <i>vlan name</i> vlan-id <i>vlan-id</i></code>                                                                                                                                                                                                                                                                    |
| Set port family to <b>ethernet-switching</b><br>Set port mode to <b>Access</b> .                                                                                                                                               | <code>set interfaces <i>interface</i> unit 0 family ethernet-switching<br/>port-mode access</code>                                                                                                                                                                                                                                |
| Set VLAN on port stanza.                                                                                                                                                                                                       | <code>set interfaces <i>interface</i> unit 0 family ethernet-switching<br/>vlan members <i>vlan-members</i></code>                                                                                                                                                                                                                |
| Set RSTP protocol with edge option.                                                                                                                                                                                            | <code>set protocols rstp interface <i>interface</i> edge</code>                                                                                                                                                                                                                                                                   |
| RSTP protocol is disabled if redundant trunk groups are configured.                                                                                                                                                            | <code>set protocols rstp interface <i>interface</i> disable</code>                                                                                                                                                                                                                                                                |
| Routed Uplink Port Role                                                                                                                                                                                                        |                                                                                                                                                                                                                                                                                                                                   |
| Set the port role to Routed Uplink.                                                                                                                                                                                            | <code>set interfaces <i>interface</i> apply-macro juniper-port-profile<br/>Routed Uplink</code>                                                                                                                                                                                                                                   |
| Set port family to inet.<br>Set IP address on the port.                                                                                                                                                                        | <code>set interfaces <i>interface</i> unit 0 family inet address<br/><i>ipaddress</i></code>                                                                                                                                                                                                                                      |
| Set class-of-service parameters<br><b>SCHEDULER_MAP=</b> <code>juniper-port-profile-map</code><br><b>IEEE_CLASSIFIER=</b> <code>juniper-ieee-classifier</code><br><b>DSCP_CLASSIFIER=</b> <code>juniper-dscp-classifier</code> | <code>set class-of-service interfaces <i>interfaces</i> scheduler-map<br/>juniper-port-profile-map<br/>set class-of-service interfaces <i>interface</i> unit 0<br/>classifiers ieee-802.1 juniper_ieee_classifier<br/>set class-of-service interfaces <i>interface</i> unit 0 classifiers<br/>dscp juniper-dscp-classifier</code> |
| Set CoS configuration                                                                                                                                                                                                          | Refer <a href="#">Table 15 on page 78</a> for details.                                                                                                                                                                                                                                                                            |
| Layer 2 Uplink Port Role                                                                                                                                                                                                       |                                                                                                                                                                                                                                                                                                                                   |
| Set the port role to <b>Layer 2 Uplink</b> .                                                                                                                                                                                   | <code>set interfaces <i>interface</i> apply-macro juniper-port-profile<br/>Layer2 Uplink</code>                                                                                                                                                                                                                                   |
| Set port family to <b>ethernet-switching</b><br>Set port mode to <b>trunk</b> .                                                                                                                                                | <code>set interfaces <i>interface</i> unit 0 family ethernet-switching<br/>port-mode trunk</code>                                                                                                                                                                                                                                 |
| Set Native VLAN name.                                                                                                                                                                                                          | <code>set interfaces <i>interface</i> unit 0 family ethernet-switching<br/>native-vlan-id <i>vlan-name</i></code>                                                                                                                                                                                                                 |
| Set the port as part of all valid VLANs; "valid" refers to all VLANs except native VLAN and voice VLANs.                                                                                                                       | <code>set interfaces <i>interface</i> unit 0 family ethernet-switching<br/>vlan members <i>vlan-members</i></code>                                                                                                                                                                                                                |
| Set port security parameter.                                                                                                                                                                                                   | <code>set ethernet-switching-options secure-access-port<br/>dhcp-trusted</code>                                                                                                                                                                                                                                                   |
| Set RSTP protocol with point-to-point option.                                                                                                                                                                                  | <code>set protocols rstp interface <i>interface</i> mode point-to-point</code>                                                                                                                                                                                                                                                    |
| Disable RSTP if redundant trunk groups are configured.                                                                                                                                                                         | <code>set protocols rstp interface <i>interface</i> disable</code>                                                                                                                                                                                                                                                                |

Table 14: Port Role Configuration Summary (*continued*)

| Configuration Description                                                                                                                                                                            | CLI Commands                                                                                                                                                                                                                                                                                                                                  |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Set class-of-service parameters.<br><br><b>SCHEDULER_MAP=</b> juniper-port-profile-map<br><br><b>IEEE_CLASSIFIER=</b> juniper_ieee_classifier<br><br><b>DSCP_CLASSIFIER=</b> juniper_dscp_classifier | <code>set class-of-service interfaces <i>interfaces</i> scheduler-map juniper-port-profile-map</code><br><code>set class-of-service interfaces <i>interface</i> unit 0 classifiers ieee-802.1 juniper_ieee_classifier</code><br><code>set class-of-service interfaces <i>interface</i> unit 0 classifiers dscp juniper-dscp-classifier</code> |
| Set CoS configuration                                                                                                                                                                                | Refer to <a href="#">Table 15 on page 78</a> for details.                                                                                                                                                                                                                                                                                     |

[Table 15 on page 78](#) lists the CLI commands for the recommended CoS settings that are committed when the CoS configuration is set.

Table 15: Recommended CoS Settings for Port Roles

| CoS Parameter                    | CLI Command                                                                                                                                                                                                                                                                                                                                                                               |
|----------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Forwarding Classes               |                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>voice</b>                     | <code>set class-of-service forwarding-classes class voice queue-num 7</code>                                                                                                                                                                                                                                                                                                              |
| <b>expedited-forwarding</b>      | <code>set class-of-service forwarding-classes class expedited-forwarding queue-num 5</code>                                                                                                                                                                                                                                                                                               |
| <b>assured-forwarding</b>        | <code>set class-of-service forwarding-classes class assured-forwarding queue-num 1</code>                                                                                                                                                                                                                                                                                                 |
| <b>best-effort</b>               | <code>set class-of-service forwarding-classes class best-effort queue-num 0</code>                                                                                                                                                                                                                                                                                                        |
| Schedulers                       |                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>strict-priority-scheduler</b> | <p>The CLI commands are:</p> <ul style="list-style-type: none"> <li><code>set class-of-service schedulers strict-priority-scheduler transmit-rate percent 10</code></li> <li><code>set class-of-service schedulers strict-priority-scheduler buffer-size percent 5</code></li> <li><code>set class-of-service schedulers strict-priority-scheduler priority strict-high</code></li> </ul> |
| <b>expedited-scheduler</b>       | <p>The CLI commands are:</p> <ul style="list-style-type: none"> <li><code>set class-of-service schedulers expedited-scheduler transmit-rate percent 30</code></li> <li><code>set class-of-service schedulers expedited-scheduler buffer-size percent 30</code></li> <li><code>set class-of-service schedulers expedited-scheduler priority low</code></li> </ul>                          |
| <b>assured-scheduler</b>         | <p>The CLI commands are:</p> <ul style="list-style-type: none"> <li><code>set class-of-service schedulers assured-scheduler transmit-rate percent 25</code></li> <li><code>set class-of-service schedulers strict-priority-scheduler buffer-size percent 25</code></li> <li><code>set class-of-service schedulers strict-priority-scheduler priority low</code></li> </ul>                |

Table 15: Recommended CoS Settings for Port Roles (*continued*)

| CoS Parameter                | CLI Command                                                                                                                                                                                                                                                                                                                 |
|------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>best-effort-scheduler</b> | <p>The CLI commands are:</p> <pre>set class-of-service schedulers best-effort-scheduler transmit-rate percent 35 set class-of-service schedulers best-effort-scheduler buffer-size percent 40 set class-of-service schedulers best-effort-scheduler priority low</pre>                                                      |
| <b>Classifiers</b>           | <p>The classifiers are:</p> <pre>set class-of-service classifiers ieee-802.1 juniper_ieee_classifier import default forwarding-class voice loss-priority low code-points 101 set class-of-service classifiers dscp juniper_dscp_classifier import default forwarding-class voice loss-priority low code-points 101110</pre> |

- Related Documentation**
- [Configuring Gigabit Ethernet Interfaces \(J-Web Procedure\) on page 68](#)
  - [Configuring Gigabit Ethernet Interfaces \(CLI Procedure\) on page 64](#)
  - [Configuring Gigabit Ethernet Interfaces \(CLI Procedure\)](#)

## Adding a Logical Unit Description to the Configuration

You can include a text description of each logical unit in the configuration file. Any descriptive text you include is displayed in the output of the **show interfaces** commands, and is also exposed in the **ifAlias** Management Information Base (MIB) object. It has no impact on the interface's configuration. To add a text description, include the **description** statement:

**description** *text*;

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

The description can be a single line of text. If the text contains spaces, enclose it in quotation marks.



**NOTE:** You can configure the extended DHCP relay to include the interface description in the option 82 Agent Circuit ID suboption. See “*Using DHCP Relay Agent Option 82 Information*” in the *Junos OS Broadband Subscriber Management and Services Library*.

For information about describing physical interfaces, see *Configuring Interface Description*.

## Disabling a Physical Interface

---

- [Disabling a Physical Interface on page 80](#)
- [Example: Disabling a Physical Interface on page 80](#)
- [Effect of Disabling Interfaces on T series PICs on page 81](#)

### Disabling a Physical Interface

You can disable a physical interface, marking it as being down, without removing the interface configuration statements from the configuration.



**CAUTION:** Dynamic subscribers and logical interfaces use physical interfaces for connection to the network. The Junos OS allows you to set the interface to disable and commit the change while dynamic subscribers and logical interfaces are still active. This action results in the loss of all subscriber connections on the interface. Use care when disabling interfaces.

To disable a physical interface:

1. In configuration mode, go to **[edit interfaces *interface-name*]** hierarchy level.

```
[edit]
user@host# edit interfaces ge-fpc/pic/port
```

2. Include the **disable** statement.

```
[edit interfaces at-fpc/pic/port]
user@host# set disable
```



**NOTE:** On the router, when you use the disable statement at the edit interfaces hierarchy level, depending on the PIC type, the interface might or might not turn off the laser. Older PIC transceivers do not support turning off the laser, but newer Gigabit Ethernet PICs with SFP and XFP transceivers do support it and the laser will be turned off when the interface is disabled.



**WARNING:** Do not stare into the laser beam or view it directly with optical instruments even if the interface has been disabled.

### Example: Disabling a Physical Interface

Sample interface configuration:

```
[edit interfaces]
user@host# show
ge-0/3/2 {
 unit 0 {
 description CE2-to-PE1;
```

```

 family inet {
 address 20.1.1.6/24;
 }
 }
}

```

Disabling the interface:

```

[edit interfaces ge-0/3/2]
user@host# set disable

```

Verifying the interface configuration:

```

[edit interfaces ge-0/3/2]
user@host# show
disable; # Interface is marked as disabled.
unit 0 {
 description CE2-to-PE1;
 family inet {
 address 20.1.1.6/24;
 }
}

```

## Effect of Disabling Interfaces on T series PICs

The following table describes the effect of using the **set interfaces disable *interface\_name*** statement on T series PICs.

**Table 16: Effect of set interfaces disable <interface\_name> on T series PICs**

| PIC Model Number  | PIC Description                                                               | Type of PIC | Behaviour             |
|-------------------|-------------------------------------------------------------------------------|-------------|-----------------------|
| PF-12XGE-SFPP     | 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (T4000 Router)                      | 5           | Tx laser disabled     |
| PF-24XGE-SFPP     | 10-Gigabit Ethernet LAN/WAN PIC with Oversubscription and SFP+ (T4000 Router) | 5           | Tx laser disabled     |
| PF-1CGE-CFP       | 100-Gigabit Ethernet PIC with CFP (T4000 Router)                              | 5           | Tx laser disabled     |
| PD-4XGE-XFP       | 10-Gigabit Ethernet, 4-port LAN/WAN XFP                                       | 4           | Tx laser disabled     |
| PD-5-10XGE-SFPP   | 10-Gigabit LAN/WAN with SFP+                                                  | 4           | Tx laser disabled     |
| PD-1XLE-CFP       | 40-Gigabit with CFP                                                           | 4           | Tx laser disabled     |
| PD-1CE-CFP-FPC4   | 100-Gigabit with CFP                                                          | 4           | Tx laser disabled     |
| PD-TUNNEL         | 40-Gigabit Tunnel Services                                                    | 4           | NA                    |
| PD-4OC192-SON-XFP | OC192/STM64, 4-port XFP                                                       | 4           | Tx laser not disabled |
| PD-1OC768-SON-SR  | OC768c/STM256, 1-port                                                         | 4           | Tx laser not disabled |

Related Documentation

- [disable on page 191](#)

## Disabling a Logical Interface

---

You can unconfigure a logical interface, effectively disabling that interface, without removing the logical interface configuration statements from the configuration. To do this, include the **disable** statement:

```
disable;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

When an interface is disabled, a route (pointing to the reserved target “**REJECT**”) with the IP address of the interface and a 32-bit subnet mask is installed in the routing table. See *Routing Protocols*.

## Configuring Flow Control

---

By default, the router or switch imposes flow control to regulate the amount of traffic sent out on a Fast Ethernet, Tri-Rate Ethernet copper, Gigabit Ethernet, and 10-Gigabit Ethernet interface. Flow control is not supported on the 4-port Fast Ethernet PIC. This is useful if the remote side of the connection is a Fast Ethernet or Gigabit Ethernet switch.

You can disable flow control if you want the router or switch to permit unrestricted traffic. To disable flow control, include the **no-flow-control** statement:

```
no-flow-control;
```

To explicitly reinstate flow control, include the **flow-control** statement:

```
flow-control;
```

You can include these statements at the following hierarchy levels:

- [edit interfaces *interface-name* aggregated-ether-options]
- [edit interfaces *interface-name* ether-options]
- [edit interfaces *interface-name* fastether-options]
- [edit interfaces *interface-name* gigether-options]



**NOTE:** On the Type 5 FPC, to prioritize control packets in case of ingress oversubscription, you must ensure that the neighboring peers support MAC flow control. If the peers do not support MAC flow control, then you must disable flow control.

---

- Related Documentation**
- [flow-control on page 201](#)
  - *Ethernet Interfaces Overview*
  - [EX Series Switches Interfaces Overview on page 3](#)
  - *Ethernet Interfaces Feature Guide for Routing Devices*

## Configuring the Interface Address

---

You assign an address to an interface by specifying the address when configuring the protocol family. For the **inet** or **inet6** family, configure the interface IP address. For the **iso** family, configure one or more addresses for the loopback interface. For the **ccc**, **ethernet-switching**, **tcc**, **mpls**, **tnp**, and **vppls** families, you never configure an address.



**NOTE:** The point-to-point (PPP) address is taken from the loopback interface address that has the primary attribute. When the loopback interface is configured as an unnumbered interface, it takes the primary address from the donor interface.

To assign an address to an interface, perform the following steps:

1. Configure the interface address at the **[edit interfaces *interface-name* unit *logical-unit-number* family *family*]** hierarchy level.
  - To configure an IPv4 address on routers and switches running Junos OS, use the **interface *interface-name* unit *number* family inet address *a.b.c.d/nn*** statement at the **[edit interfaces]** hierarchy level.

```
[edit interfaces]
```

```
user@host# set interface-name unit logical-unit-number family inet address a.b.c.d/nn
```



#### NOTE:

- Juniper Networks routers and switches support /31 destination prefixes when used in point-to-point Ethernet configurations; however, they are not supported by many other devices, such as hosts, hubs, routers, or switches. You must determine if the peer system also supports /31 destination prefixes before configuration.
- You can configure the same IPv4 address on multiple physical interfaces. When you assign the same IPv4 address to multiple physical interfaces, the operational behavior of those interfaces differs, depending on whether they are implicitly or explicitly point-to-point .
- By default, all interfaces are assumed to be point-to-point (PPP) interfaces. For all interfaces except aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet, you can explicitly configure an interface to be a point-to-point connection.
- If you configure the same IP address on multiple interfaces in the same routing instance, Junos OS uses only the first configuration. The remaining IP address configurations are ignored, leaving some interfaces without an assigned address. Interfaces without an assigned address cannot be used as a donor interface for an unnumbered Ethernet interface.

- To configure an IPv6 address on routers and switches running Junos OS, use the **interface *interface-name* unit *number* family inet6 address *aaaa:bbbb:::zzzz/nn*** statement at the **[edit interfaces]** hierarchy level.

```
[edit interfaces]
```

```
user@host# set interface-name unit logical-unit-number family inet6 address
aaaa:bbbb:::zzzz/nn
```



#### NOTE:

- You represent IP version 6 (IPv6) addresses in hexadecimal notation using a colon-separated list of 16-bit values. The double colon (::) represents all bits set to 0.
- You must manually configure the router or switch advertisement and advertise the default prefix for autoconfiguration to work on a specific interface.

2. [Optional] Set the broadcast address on the network or subnet .

```
[edit interfaces interface-name unit logical-unit-number family family address address],
user@host# set broadcast address
```



**NOTE:** The broadcast address must have a host portion of either all ones or all zeros. You cannot specify the addresses 0.0.0.0 or 255.255.255.255

3. [Optional] specify the remote address of the connection for the encrypted, PPP-encapsulated, and tunnel interfaces.

```
[edit logical-systems logical-system-name interfaces interface-name unit
 logical-unit-number family family address address]
user@host# set destination address
```

4. [Optional] For interfaces that carry IP version 6 (IPv6) traffic, configure the host to assign itself a unique 64-Bit IP Version 6 interface identifier (EUI-64).

```
[edit logical-systems logical-system-name interfaces interface-name unit
 logical-unit-number family family address address]
user@host# set eui-64
```

#### Related Documentation

- [Configuring Default, Primary, and Preferred Addresses and Interfaces](#)

## Configuring the Interface Bandwidth

By default, the Junos OS uses the physical interface's speed for the MIB-II object, **ifSpeed**. You can configure the logical unit to populate the **ifSpeed** variable by configuring a bandwidth value for the logical interface. The **bandwidth** statement sets an informational-only parameter; you cannot adjust the actual bandwidth of an interface with this statement.



**NOTE:** We recommend that you be careful when setting this value. Any interface bandwidth value that you configure using the **bandwidth** statement affects how the interface cost is calculated for a dynamic routing protocol, such as OSPF. By default, the interface cost for a dynamic routing protocol is calculated using the following formula:

$$\text{cost} = \text{reference-bandwidth} / \text{bandwidth},$$

where bandwidth is the physical interface speed. However, if you specify a value for bandwidth using the **bandwidth** statement, that value is used to calculate the interface cost, rather than the actual physical interface bandwidth.

To configure the bandwidth value for a logical interface, include the **bandwidth** statement:

```
bandwidth rate;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

**rate** is the peak rate, in bps or cps. You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000). You can also specify a value in cells per second by entering a decimal number followed by the abbreviation **c**; values expressed in cells per second are converted to bits per second using the formula 1 cps = 384 bps. The value can be any positive integer. The **bandwidth** statement is valid for all logical interfaces, except multilink interfaces.

## Configuring the Media MTU

The media maximum transmission unit (MTU) is the largest data unit that can be forwarded without fragmentation. The default media MTU size used on a physical interface depends on the encapsulation being used on that interface. For a listing of MTU sizes for each encapsulation type, see *Media MTU Sizes by Interface Type*.

To configure the media-MTU size:

1. In configuration mode, go to the [edit interfaces *interface-name*] hierarchy level.

```
[edit]
user@host# [edit interfaces interface-name]
```

2. Include the **mtu** statement.

```
[edit interfaces interface-name]
mtu bytes;
```

- If you change the size of the media MTU, you must ensure that the size is equal to or greater than the sum of the protocol MTU and the encapsulation overhead. You configure the protocol MTU by including the **mtu** statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family *family*]

- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family *family*]



## NOTE:

- Changing the media MTU or protocol MTU causes an interface to be deleted and added again.
- Because tunnel services interfaces are considered logical interfaces, you cannot configure the MTU setting for the physical interface. This means you cannot include the `mtu` statement at the [edit interfaces *interface-name*] hierarchy level for the following interface types: generic routing encapsulation (*gr-*), IP-IP (*ip-*), loopback (*lo-*), link services (*ls-*), multilink services (*ml-*), and multicast (*pe-*, *pd-*). You can, however, configure the protocol MTU on tunnel interfaces.
- If you configure an MTU value by including the `mtu` statement at the [edit interfaces *interface-name* unit *logical-unit-number* family *mpls*] hierarchy level, the configured value is used.

## Related Documentation

- [Media MTU Overview](#)
- [Media MTU Sizes by Interface Type](#)
- [Encapsulation Overhead by Interface Encapsulation Type](#)
- [Setting the Protocol MTU on page 88](#)

## Setting the Protocol MTU

When you initially configure an interface, the protocol maximum transmission unit (MTU) is calculated automatically. If you subsequently change the media MTU, the protocol MTU on existing address families automatically changes.

For a list of default protocol MTU values, see [Media MTU Sizes by Interface Type](#).

To modify the MTU for a particular protocol family, include the `mtu` statement:

`mtu bytes;`

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family *family*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family *family*]

If you increase the size of the protocol MTU, you must ensure that the size of the media MTU is equal to or greater than the sum of the protocol MTU and the encapsulation overhead. For a list of encapsulation overhead values, see [Encapsulation Overhead by Interface Encapsulation Type](#). If you reduce the media MTU size, but there are already one or more address families configured and active on the interface, you must also reduce

the protocol MTU size. (You configure the media MTU by including the **mtu** statement at the **[edit interfaces *interface-name*]** hierarchy level.)



**NOTE:** Changing the media MTU or protocol MTU causes an interface to be deleted and added again.

The maximum number of data-link connection identifiers (DLCIs) is determined by the MTU on the interface. If you have keepalives enabled, the maximum number of DLCIs is 1000, with the MTU set to 5012.

The actual frames transmitted also contain cyclic redundancy check (CRC) bits, which are not part of the MTU. For example, the default protocol MTU for a Gigabit Ethernet interface is 1500 bytes, but the largest possible frame size is actually 1504 bytes; you need to consider the extra bits in calculations of MTUs for interoperability.

#### Related Documentation

- [Media MTU Overview](#)
- [Configuring the Media MTU on page 87](#)

## Interface Ranges



**NOTE:** This task uses Junos OS for EX Series switches that does not support the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that supports ELS, see *Configuring Interface Ranges*. For ELS details, see *Getting Started with Enhanced Layer 2 Software*.

The Junos OS allows you to group a range of identical interfaces into an *interface range*. You first specify the group of identical interfaces in the interface range. Then you can apply a common configuration to the specified interface range, reducing the number of configuration statements required and saving time while producing a compact configuration.

- [Configuring Interface Ranges on page 89](#)
- [Expanding Interface Range Member and Member Range Statements on page 93](#)
- [Configuration Inheritance for Member Interfaces on page 94](#)
- [Member Interfaces Inheriting Configuration from Configuration Groups on page 95](#)
- [Interfaces Inheriting Common Configuration on page 96](#)
- [Configuring Inheritance Range Priorities on page 97](#)
- [Configuration Expansion Where Interface Range Is Used on page 97](#)

## Configuring Interface Ranges

To configure an interface range, include the **interface-range** statement at the **[edit interfaces]** hierarchy level.

The **interface-range** statement accepts only physical networking interface names in its definition. The following interface types are supported and example CLI descriptors are shown:

- ATM—**at-fpc/pic/port**
- Channelized—(**coc | cstm**)**n-fpc/pic/port**
- DPC—**xe-fpc/pic/port**
- E1/E3—(**e1 | e3**)-**fpc/pic/port**
- Ethernet—(**xe | ge | fe**)-**fpc/pic/port**
- ISDN—**isdn-fpc/pic/port**
- Serial—**se-fpc/pic/port**
- SONET/SDH—**so-fpc/pic/port**
- T1/T3—(**t1 | t3**)-**fpc/pic/port**

Interfaces can be grouped either as a range of interfaces or using a number range under the **interface-range** statement definition.

Interfaces in an **interface-range** definition can be added as part of a member range or as individual members or multiple members using a number range.

To specify a member range, use the **member-range** statement at the **[edit interfaces interface-range name]** hierarchy level.

To specify interfaces in lexical order, use the **member-range start-range to end-range** statement.

A range for a member statement should contain the following:

- **\***—All, specifies sequential interfaces from 0 through 47.



**CAUTION:** The wildcard **\*** in a member statement does not take into account the interface numbers supported by a specific interface type. Irrespective of the interface type, **\*** includes interface numbers ranging from 0 through 47 to the interface group. Therefore, use **\*** in a member statement with caution.

- **num**—Number, specifies one specific interface by its number.
- **[low-high]**—Numbers between low to high, specifies a range of sequential interfaces.
- **[num1, num2, num3]**—Numbers **num1**, **num2**, and **num3** specify multiple specific interfaces.

**Example: Specifying an  
Interface Range  
Member Range**

```
member-range ge-0/0/0 to ge-4/0/40;
```

To specify one or multiple members, use the **member** statement at the **[edit interfaces interface-range name]** hierarchy level.

To specify the list of interface range members individually or for multiple interfaces using regex, use the **member list of interface names** statement.

**Example: Specifying an Interface Range Member**

```
member ge-0/0/0;
member ge-0/*/*
member ge-0/[1-10]/0;
member ge-0/[1,2,3]/3;
```

Regex or wildcards are not supported for interface-type prefixes. For example, prefixes **ge**, **fe**, and **xe** must be mentioned explicitly.

An **interface-range** definition can contain both **member** and **member-range** statements within it. There is no maximum limit on the number of **member** or **member-range** statements within an interface-range. However, at least one **member** or **member-range** statement must exist within an **interface-range** definition.

**Example: Interface Range Common Configuration**

Configuration common to an interface range can be added as a part of the **interface-range** definition, as follows:

```
[edit]
interfaces {
 + interface-range foo {
 + member-range ge-1/0/0 to ge-4/0/40;
 + member ge-0/1/1;
 + member ge-5/[1-10]/*;
 /*Common configuration is added as part of interface-range definition*/
 mtu 256;
 hold-time up 10;
 ether-options {
 flow-control;
 speed {
 100m;
 }
 802.3ad primary;
 }
 }
}
```

An **interface-range** definition having just **member** or **member-range** statements and no common configurations statements is valid.

These defined interface ranges can be used in other configuration hierarchies, in places where an **interface** node exists.

**Example: Interface-Range foo Used Under the Protocols Hierarchy**

```
protocols {
 dot1x {
 authenticator {
 interface foo{
 retries 1;
 }
 }
 }
}
```

**foo** should be an **interface-range** defined at the **[interfaces]** hierarchy level. In the above example, the **interface** node can accept both individual interfaces and interface ranges.



**TIP:** To view an interface range in expanded configuration, use the **(show | display inheritance)** command. For more information, see the *CLI User Guide*.

By default, **interface-range** is not available to configure in the CLI where the **interface** statement is available. The following locations are supported; however, some of the hierarchies shown in this list are product specific:

- protocols dot1x authentication interface
- protocols dvmrp interface
- protocols oam ethernet lmi interface
- protocols esis interface
- protocols igmp interface
- protocols igmp-host client *num* interface
- protocols mld-host client *num* interface
- protocols router-advertisement interface
- protocols isis interface
- protocols ldp interface
- protocols oam ethernet link-fault-management interface
- protocols lldp interface
- protocols link-management peer lmp-control-channel interface
- protocols link-management peer control-channel
- protocols link-management te-link *name* interface
- protocols mld interface
- protocols ospf area *id* interface
- protocols pim interface
- protocols router-discovery interface
- protocols rip group *name* neighbour
- protocols ripng group *name* neighbour
- protocols rsvp interface
- protocols snmp interface
- protocols layer2-control bpdu-block interface
- protocols layer2-control mac-rewrite interface

- protocols mpls interface
- protocols stp interface
- protocols rstp interface
- protocols mstp interface
- protocols vstp interface
- protocols mstp msti *id* interface
- protocols mstp msti vlan *id* interface
- protocols vstp vlan *name* interface
- protocols gvrp interface
- protocols igmp-snooping vlan *name* interface
- protocols lldp interface
- protocols lldp-med interface
- protocols sflow interfaces
- ethernet-switching-options analyzer *name* input [egress | ingress ] interface
- ethernet-switching-options analyzer *name* output interface
- ethernet-switching-options secure-access-port interface
- ethernet-switching-options interfaces ethernet-switching-options voip interface
- ethernet-switching-options redundant-trunk-group group *g1* interface
- ethernet-switching-options redundant-trunk-group group *g1* interface
- ethernet-switching-options bpdu-block interface
- poe interface vlans pro-bng-mc1-bsd1 interface

## Expanding Interface Range Member and Member Range Statements

All **member** and **member-range** statements in an interface range definition are expanded to generate the final list of interface names for the specified interface range.

### Example: Expanding Interface Range Member and Member Range Statements

```
[edit]
interfaces {
 interface-range range-1 {
 member-range ge-0/0/0 to ge-4/0/20;
 member ge-10/1/1;
 member ge-5/[0-5]/*;
 /*Common configuration is added part of the interface-range definition*/
 mtu 256;
 hold-time up 10;
 ether-options {
 flow-control;
 speed {
 100m;
 }
 }
 }
}
```

```

 802.3ad primary;
 }
}

```

For the **member-range** statement, all possible interfaces between **start-range** and **end-range** are considered in expanding the members. For example, the following **member-range** statement:

```
member-range ge-0/0/0 to ge-4/0/20
```

expands to:

```

[ge-0/0/0, ge-0/0/1 ... ge-0/0/max_ports
 ge-0/1/0 ge-0/1/1 ... ge-0/1/max_ports
 ge-0/2/0 ge-0/2/1 ... ge-0/2/max_ports
 .
 .
 ge-0/MAX_PICS/0 ... ge-0/max_pics/max_ports
 ge-1/0/0 ge-1/0/1 ... ge-1/0/max_ports
 .
 ge-1/MAX_PICS/0 ... ge-1/max_pics/max_ports
 .
 .
 ge-4/0/0 ge-4/0/1 ... ge-4/0/max_ports]

```

The following **member** statement:

```
ge-5/[0-5]/*
```

expands to:

```

ge-5/0/0 ... ge-5/0/max_ports
ge-5/1/0 ... ge-5/0/max_ports
.
.
ge-5/5/0 ... ge-5/5/max_ports

```

The following **member** statement:

```
ge-5/1/[2,3,6,10]
```

expands to:

```

ge-5/1/2
ge-5/1/3
ge-5/1/6
ge-5/1/10

```

## Configuration Inheritance for Member Interfaces

When the Junos OS expands the **member** and **member-range** statements present in an **interface-range**, it creates *interface objects* if they are not explicitly defined in the configuration. The common configuration is copied to all its member interfaces in the **interface-range**.

**Example:** Foreground interface configuration takes priority compared to configuration inherited by the interface through the **interface-range**.

### Configuration Priorities

```
interfaces {
```

```

interface-range range-1 {
 member-range ge-1/0/0/ to ge-10/0/47;
 mtu 256;
}
ge-1/0/1 {
 mtu 1024;
}
}

```

In the preceding example, interface **ge-1/0/1** will have an MTU value of 1024.

This can be verified with output of the **show interfaces | display inheritance** command, as follows:

```

user@host: # show interfaces | display inheritance
'ge-1/0/0' was expanded from interface-range 'range-1'
##
ge-1/0/0 {
 ##
 ## '256' was expanded from interface-range 'range-1'
 ##
 mtu 256;
}
ge-1/0/1 {
 mtu 1024;
}
##
'ge-1/0/2' was expanded from interface-range 'range-1'
##
ge-1/0/2 {
 ##
 ## '256' was expanded from interface-range 'range-1'
 ##
 mtu 256;
}
.....
.....
##
'ge-10/0/47' was expanded from interface-range 'range-1'
##
ge-10/0/47 {
 ##
 ## '256' was expanded from interface-range 'range-1'
 ##
 mtu 256;
}

```

## Member Interfaces Inheriting Configuration from Configuration Groups

Interface range member interfaces inherit the config-groups configuration like any other foreground configuration. **interface-range** is similar to any other foreground configuration statement. The only difference is that the **interface-range** goes through a member interfaces expansion before Junos OS reads this configuration.

```

groups {
 global {
 interfaces {
 <*> {
 hold-time up 10;

```

```

 }
 }
}
apply-groups [global];
interfaces {
 interface-range range-1 {
 member-range ge-1/0/0 to ge-10/0/47;
 mtu 256;
 }
}
}

```

The **hold-time** configuration is applied to all members of **interface-range range-1**.

This can be verified with **show interfaces | display inheritance** as follows:

```

user@host# show interfaces | display inheritance
ge-1/0/0 {
 ##
 ## '256' was expanded from interface-range 'range-1'
 ##
 mtu 256;
 ##
 ## 'hold-time' was inherited from group 'global'
 ## '10' was inherited from group 'global'
 ##
 hold-time up 10;
}
ge-1/0/1 {
 ##
 ## '256' was expanded from interface-range 'range-1'
 ##
 mtu 256;
 ##
 ## 'hold-time' was inherited from group 'global'
 ## '10' was inherited from group 'global'
 ##
 hold-time up 10;
}
ge-10/0/47 {
 ##
 ## '256' was expanded from interface-range 'range-1'
 ##
 mtu 256;
 ##
 ## 'hold-time' was inherited from group 'global'
 ## '10' was inherited from group 'global'
 ##
 hold-time up 10;
}

```

## Interfaces Inheriting Common Configuration

If an interface is a member of several interface ranges, that interface will inherit the common configuration from all of those interface ranges.

```

[edit]
interfaces {
 interface-range range-1 {

```

```

 member-range ge-1/0/0 to ge-10/0/47;
 mtu 256;
 }
}
interfaces {
 interface-range range-1 {
 member-range ge-10/0/0 to ge-10/0/47;
 hold-time up 10;
 }
}

```

In this example, interfaces **ge-10/0/0** through **ge-10/0/47** will have both **hold-time** and **mtu**.

## Configuring Inheritance Range Priorities

The interface ranges are defined in the order of inheritance priority, with the first interface range configuration data taking priority over subsequent interface ranges.

```

[edit]
interfaces {
 interface-range int-grp-one {
 member-range ge-0/0/0 to ge-4/0/40;
 member ge-1/1/1;
 /*Common config is added part of the interface-range definition*/
 mtu 256;
 hold-time up 10;
 }
}
interfaces {
 interface-range int-grp-two {
 member-range ge-5/0/0 to ge-10/0/40;
 member ge-1/1/1;
 mtu 1024;
 }
}

```

Interface **ge-1/1/1** exists in both **interface-range int-grp-one** and **interface-range int-grp-two**. This interface inherits **mtu 256** from **interface-range int-grp-one** because it was defined first.

## Configuration Expansion Where Interface Range Is Used

In this example, **interface-range range-1** is used under the **protocols** hierarchy:

```

[edit]
interfaces {
 interface-range range-1 {
 member ge-10/1/1;
 member ge-5/5/1;
 mtu 256;
 hold-time up 10;
 ether-options {
 flow-control;
 speed {
 100m;
 }
 }
 }
}

```

```
 802.3ad primary;
 }
}
protocols {
 dot1x {
 authenticator {
 interface range-1 {
 retries 1;
 }
 }
 }
}
}
```

The **interface** node present under **authenticator** is expanded into member interfaces of the **interface-range** *range-1* as follows:

```
protocols {
 dot1x {
 authenticator {
 interface ge-10/1/1 {
 retries 1;
 }
 interface ge-5/5/1 {
 retries 1;
 }
 }
 }
}
```

The **interface** *range-1* statement is expanded into two interfaces, ge-10/1/1 and ge-5/5/1, and configuration **retries 1** is copied under those two interfaces.

This configuration can be verified using the **show protocols dot1x | display inheritance** command.

**Related Documentation**

- [Physical Interfaces](#)

---

## Configuring Accounting for the Physical Interface

- [Accounting Profiles Overview on page 98](#)
- [Configuring Accounting for the Physical Interface on page 99](#)
- [Displaying Accounting Profile for the Physical Interface on page 100](#)

### Accounting Profiles Overview

Juniper Networks routers and switches can collect various kinds of data about traffic passing through the router and switch. You can set up one or more *accounting profiles* that specify some common characteristics of this data, including the following:

- The fields used in the accounting records
- The number of files that the router or switch retains before discarding, and the number of bytes per file

- The polling period that the system uses to record the data

You configure the profiles and define a unique name for each profile using statements at the **[edit accounting-options]** hierarchy level. There are two types of accounting profiles: interface profiles and filter profiles. You configure interface profiles by including the **interface-profile** statement at the **[edit accounting-options]** hierarchy level. You configure filter profiles by including the **filter-profile** statement at the **[edit accounting-options]** hierarchy level. For more information, see the *Network Management Administration Guide for Routing Devices*.

You apply filter profiles by including the **accounting-profile** statement at the **[edit firewall filter *filter-name*]** and **[edit firewall family *family* filter *filter-name*]** hierarchy levels. For more information, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*.

## Configuring Accounting for the Physical Interface

### Before you begin

You must configure a profile to collect error and statistic information for input and output packets on a particular physical interface. An accounting profile specifies what statistics should be collected and written to a log file. For more information on how to configure an accounting-data log file, see the *Configuring Accounting-Data Log Files*.

An interface profile specifies the information collected and written to a log file. You can configure a profile to collect error and statistic information for input and output packets on a particular physical interface.

1. To configure which statistics should be collected for an interface, include the **fields** statement at the **[edit accounting-options interface-profile *profile-name*]** hierarchy level.

```
[edit accounting-options interface-profile profile-name]
user@host# set fields field-name
```

2. Each accounting profile logs its statistics to a file in the **/var/log** directory. To configure which file to use, include the **file** statement at the **[edit accounting-options interface-profile *profile-name*]** hierarchy level.

```
[edit accounting-options interface-profile profile-name]
user@host# set file filename
```



**NOTE:** You must specify a **file** statement for the interface profile that has already been configured at the **[edit accounting-options]** hierarchy level. For more information, see the [Configuring Accounting-Data Log Files](#)

3. Each interface with an accounting profile enabled has statistics collected once per interval time specified for the accounting profile. Statistics collection time is scheduled evenly over the configured interval. To configure the interval, include the **interval** statement at the **[edit accounting-options interface-profile *profile-name*]** hierarchy level.

```
[edit accounting-options interface-profile profile-name]
user@host# set interval minutes
```



**NOTE:** The minimum interval allowed is 1 minute. Configuring a low interval in an accounting profile for a large number of interfaces might cause serious performance degradation.

4. To configure the interfaces on which the accounting needs to be performed, apply the interface profile to a physical interface by including the **accounting-profile** statement at the **[edit interfaces interface-name]** hierarchy level.

```
[edit interfaces]
user@host# set interface-name accounting-profile profile-name
```

## Displaying Accounting Profile for the Physical Interface

**Purpose** To display the configured accounting profile a particular physical interface at the **[edit accounting-options interface-profile *profile-name*]** hierarchy level:

- interface-name—ge-1/0/1
- Interface profile —if\_profile
- File name—if\_stats
- Interval—15 minutes

**Action** • Run the **show** command at the **[edit edit interfaces ge-1/0/1]** hierarchy level.

```
[edit interfaces ge-1/0/1]
accounting-profile if_profile;
```

- Run the **show** command at the **[edit accounting-options]** hierarchy level.

```
interface-profile if_profile {
 interval 15;
 file if_stats {
 fields {
 input-bytes;
 output-bytes;
 input-packets;
 output-packets;
 input-errors;
 output-errors;
 }
 }
}
```

**Meaning** The configured accounting and its associated set options are displayed as expected.

## Configuring Accounting for the Logical Interface

- [Accounting Profiles Overview on page 101](#)
- [Configuring Accounting for the Logical Interface on page 101](#)
- [Displaying Accounting Profile for the Logical Interface on page 102](#)

### Accounting Profiles Overview

Juniper Networks routers and switches can collect various kinds of data about traffic passing through the router and switch. You can set up one or more *accounting profiles* that specify some common characteristics of this data, including the following:

- The fields used in the accounting records
- The number of files that the router or switch retains before discarding, and the number of bytes per file
- The polling period that the system uses to record the data

You configure the profiles and define a unique name for each profile using statements at the **[edit accounting-options]** hierarchy level. There are two types of accounting profiles: interface profiles and filter profiles. You configure interface profiles by including the **interface-profile** statement at the **[edit accounting-options]** hierarchy level. You configure filter profiles by including the **filter-profile** statement at the **[edit accounting-options]** hierarchy level. For more information, see the *Network Management Administration Guide for Routing Devices*.

You apply filter profiles by including the **accounting-profile** statement at the **[edit firewall filter *filter-name*]** and **[edit firewall family *family* filter *filter-name*]** hierarchy levels. For more information, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*.

## Configuring Accounting for the Logical Interface

### Before you begin

You must configure a profile to collect error and statistic information for input and output packets on a particular logical interface. An accounting profile specifies what statistics should be collected and written to a log file. For more information on how to configure an accounting-data log file, see the *Configuring Accounting-Data Log Files*.

An interface profile specifies the information collected and written to a log file. You can configure a profile to collect error and statistic information for input and output packets on a particular logical interface.

1. To configure which statistics should be collected for an interface, include the **fields** statement at the **[edit accounting-options interface-profile *profile-name*]** hierarchy level.

```
[edit accounting-options interface-profile profile-name]
user@host# set fields field-name
```

- Each accounting profile logs its statistics to a file in the `/var/log` directory. To configure which file to use, include the `file` statement at the `[edit accounting-options interface-profile profile-name]` hierarchy level.

```
[edit accounting-options interface-profile profile-name]
user@host# set file filename
```



**NOTE:** You must specify a file statement for the interface profile that has already been configured at the `[edit accounting-options]` hierarchy level. For more information, see the [Configuring Accounting-Data Log Files](#)

- Each interface with an accounting profile enabled has statistics collected once per interval time specified for the accounting profile. Statistics collection time is scheduled evenly over the configured interval. To configure the interval, include the `interval` statement at the `[edit accounting-options interface-profile profile-name]` hierarchy level.

```
[edit accounting-options interface-profile profile-name]
user@host# set interval minutes
```



**NOTE:** The minimum interval allowed is 1 minute. Configuring a low interval in an accounting profile for a large number of interfaces might cause serious performance degradation.

- To configure the interfaces on which the accounting needs to be performed, apply the interface profile to a logical interface by including the `accounting-profile` statement at the `[edit interfaces interface-name unit logical-unit-number]` hierarchy level.

```
[edit interfaces]
user@host# set interface-name unit logical-unit-number accounting-profile profile-name
```

## Displaying Accounting Profile for the Logical Interface

**Purpose** To display the configured accounting profile a particular logical interface at the `[edit accounting-options interface-profile profile-name]` hierarchy level:

- interface-name—`ge-1/0/1`
- Logical unit number—`1`
- Interface profile —`if_profile`
- File name—`if_stats`
- Interval—15 minutes

**Action** • Run the `show` command at the `[edit interfaces ge-1/0/1 unit 1]` hierarchy level.

```
[edit interfaces ge-1/0/1 unit 1]
accounting-profile if_profile;
```

- Run the `show` command at the `[edit accounting-options]` hierarchy level.

```

interface-profile if_profile {
 interval 15;
 file if_stats {
 fields {
 input-bytes;
 output-bytes;
 input-packets;
 output-packets;
 input-errors;
 output-errors;
 }
 }
}

```

**Meaning** The configured accounting and its associated set options are displayed as expected.

## Configuring Ethernet Loopback Capability

By default, local aggregated Ethernet, Fast Ethernet, Tri-Rate Ethernet copper, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces connect to a remote system. To place an interface in loopback mode, include the **loopback** statement:

**loopback;**



**NOTE:** If you configure a local loopback on a 1-port 10-Gigabit IQ2 and IQ2-E PIC using the **loopback** statement at the [edit interfaces *interface-name* *gigether-options*] hierarchy level, the transmit-path stops working, causing the remote end to detect a link down.

To return to the default—that is, to disable loopback mode—delete the **loopback** statement from the configuration:

```

[edit]
user@host# delete interfaces fe-fpc/pic/port fastether-options loopback

```

To explicitly disable loopback mode, include the **no-loopback** statement:

**no-loopback;**

You can include the **loopback** and **no-loopback** statements at the following hierarchy levels:

- [edit interfaces *interface-name* aggregated-ether-options]
- [edit interfaces *interface-name* ether-options]
- [edit interfaces *interface-name* fastether-options]
- [edit interfaces *interface-name* gigether-options]

**Related Documentation**

- [loopback on page 224](#)
- [Ethernet Interfaces Overview](#)

- [EX Series Switches Interfaces Overview on page 3](#)
- *Ethernet Interfaces Feature Guide for Routing Devices*

---

## Configuring Gratuitous ARP

Gratuitous Address Resolution Protocol (ARP) requests help detect duplicate IP addresses. A gratuitous ARP is a broadcast request for a router's own IP address. If a router or switch sends an ARP request for its own IP address and no ARP replies are received, the router- or switch-assigned IP address is not being used by other nodes. However, if a router or switch sends an ARP request for its own IP address and an ARP reply is received, the router- or switch-assigned IP address is already being used by another node.

Gratuitous ARP replies are reply packets sent to the broadcast MAC address with the target IP address set to be the same as the sender's IP address. When the router or switch receives a gratuitous ARP reply, the router or switch can insert an entry for that reply in the ARP cache. By default, updating the ARP cache on gratuitous ARP replies is disabled on the router or switch.

To enable updating of the ARP cache for gratuitous ARPs:

1. In configuration mode, go to the **[edit interfaces interface-name]** hierarchy level.

```
[edit]
user@host# edit interfaces interface-name
```

2. Include the **gratuitous-arp-reply** statement.

```
[edit interfaces interface-name]
user@host# set gratuitous-arp-reply
```

To restore the default behavior, that is, to disable updating of the ARP cache for gratuitous ARP, delete the **gratuitous-arp-reply** statement from the configuration:

```
[edit interfaces interface-name]
user@host# delete gratuitous-arp-reply;
```

By default, the router or switch responds to gratuitous ARP requests. However, on Ethernet interfaces, you can disable responses to gratuitous ARP requests.

To disable responses to gratuitous ARP requests:

1. In configuration mode, go to the **[edit interfaces interface-name]** hierarchy level.

```
[edit]
user@host# edit interfaces interface-name
```

2. Include the **no-gratuitous-arp-request** statement.

```
[edit interfaces interface-name]
user@host# set no-gratuitous-arp-request
```

To return to the default—that is, to respond to gratuitous ARP requests—delete the **no-gratuitous-arp-request** statement from the configuration:

```
[edit interfaces interface-name]
user@host# delete no-gratuitous-arp-request
```

#### Related Documentation

- [gratuitous-arp-reply on page 202](#)
- [no-gratuitous-arp-request](#)
- [Ethernet Interfaces Overview](#)
- [EX Series Switches Interfaces Overview on page 3](#)
- [Ethernet Interfaces Feature Guide for Routing Devices](#)

## Configuring Static ARP Table Entries For Mapping IP Addresses to MAC Addresses

By default, the device responds to an Address Resolution Protocol (ARP) request only if the destination address of the ARP request is on the local network of the incoming interface. For Fast Ethernet or Gigabit Ethernet interfaces, you can configure static ARP entries that associate the IP addresses of nodes on the same Ethernet subnet with their media access control (MAC) addresses. These static ARP entries enable the device to respond to ARP requests even if the destination address of the ARP request is not local to the incoming Ethernet interface.

Also, unlike dynamically learned ARP entries, static ARP entries do not age out. You can also configure static ARP entries in a troubleshooting situation or if your device is unable to learn a MAC address dynamically.



**NOTE:** By default, an ARP policer is installed that is shared among all the Ethernet interfaces on which you have configured the `family inet` statement. By including the `arp` statement at the `[edit interfaces interface-name unit logical-unit-number family inet policer]` hierarchy level, you can apply a specific ARP-packet policer to an interface. This feature is not available on EX Series switches.

To configure static ARP entries:

1. In the configuration mode, at the `[edit]` hierarchy level, configure the router interface on which the ARP table entries for the router is configured.

```
[edit]
user@host# edit interfaces interface-name
```

2. Configure the protocol family, the logical unit of the interface, and the interface address of the router interface at the `[edit interfaces interface-name]` hierarchy level. While configuring the protocol family, specify `inet` as the protocol family.



**NOTE:** When you need to conserve IP addresses, you can configure an Ethernet interface to be unnumbered by including the `unnumbered-address` statement at the `[edit interfaces interface-name unit logical-unit-number family inet]` hierarchy level.

```
[edit interfaces interface-name]
```

```
user@host# edit unit logical-unit-number family inet address interface-address
```

3. Configure a static ARP entry by specifying the IP address and the MAC address that are to be mapped to each other. The IP address specified must be part of the subnet defined in the enclosing **address** statement. The MAC address must be specified as hexadecimal bytes in the following formats: *nnnn.nnnn.nnnn* or *nn:nn:nn:nn:nn:nn* format. For instance, you can use either **0011.2233.4455** or **00:11:22:33:44:55**.

```
[edit interfaces interface-name unit logical-unit-number family inet address
 interface-address
```

```
user@host# set arp ip-address mac mac-address
```

4. Configure another static ARP entry by specifying the IP address and the MAC address that are to be mapped to each other. You can also associate a multicast MAC address with a unicast IP address by including the **multicast-mac** option with the **arp** statement. You can optionally configure the router to respond to ARP requests for the specified IP address by using the **publish** option with the **arp** statement.



**NOTE:** For unicast MAC addresses only, if you include the **publish** option, the router or switch replies to proxy ARP requests.

```
[edit interfaces interface-name unit logical-unit-number family inet address
 interface-address
```

```
user@host# set arp ip-address multicast-mac mac-address publish
```



**NOTE:** The Junos OS supports the IPv6 static neighbor discovery cache entries, similar to the static ARP entries in IPv4.

#### Related Documentation

- [arp on page 182](#)
- *Static ARP Table Entries Overview*
- *Management Ethernet Interface Overview*
- [EX Series Switches Interfaces Overview on page 3](#)
- *Applying Policers*
- *Configuring an Unnumbered Interface*
- *Ethernet Interfaces Feature Guide for Routing Devices*

## Disabling the Transmission of Redirect Messages on an Interface

By default, the interface sends protocol redirect messages. To disable the sending of these messages on an interface, include the **no-redirects** statement:

```
no-redirects;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family *family*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family *family*]

To disable the sending of protocol redirect messages for the entire router or switch, include the **no-redirects** statement at the [edit system] hierarchy level.

Related  
Documentation

- [no-redirects on page 239](#)

## Configuring Restricted and Unrestricted Proxy ARP

To configure restricted or unrestricted proxy ARP, include the **proxy-arp** statement:

**proxy-arp** (restricted |unrestricted);

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* ]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

To return to the default—that is, to disable restricted or unrestricted proxy ARP—delete the **proxy-arp** statement from the configuration:

```
[edit]
user@host# delete interfaces interface-name unit logical-unit-number proxy-arp
```

You can track the number of restricted or unrestricted proxy ARP requests processed by the router or switch by issuing the **show system statistics arp** operational mode command.



**NOTE:** When proxy ARP is enabled as default or unrestricted, the router or switch responds to any ARP request as long as the device has an active route to the target address of the ARP request. This gratuitous ARP behavior can result in an error when the receiving interface and target response interface are the same and the end device (for example, a client) performs a duplicate address check. To prevent this error, configure the router or switch interface with the **no-gratuitous-arp-reply** statement. See “[Configuring Gratuitous ARP on page 104](#)” for information about how to disable responses to gratuitous ARP requests.

Related  
Documentation

- [proxy-arp on page 249](#)
- *Restricted and Unrestricted Proxy ARP Overview*
- [Configuring Gratuitous ARP on page 104](#)
- *Ethernet Interfaces Feature Guide for Routing Devices*

## Enabling or Disabling SNMP Notifications on Logical Interfaces

---

By default, Simple Network Management Protocol (SNMP) notifications are sent when the state of an interface or a connection changes. To explicitly enable these notifications on the logical interface, include the **traps** statement; to disable these notifications on the logical interface, include the **no-traps** statement:

```
(traps | no-traps);
```

You can include these statements at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

## Enabling or Disabling SNMP Notifications on Physical Interfaces

---

By default, Simple Network Management Protocol (SNMP) notifications are sent when the state of an interface or a connection changes. You can enable or disable these notification based on you requirements.

To explicitly enable sending SNMP notifications on the physical interface, perform the following steps:

1. In configuration mode, go to the [edit interfaces *interface-name*] hierarchy level:

```
[edit]
user@host# edit interfaces interface-name
```

2. Configure the **traps** option to enable sending of Simple Network Management Protocol (SNMP) notifications when the state of the connection changes.

```
[edit interfaces interface-name]
user@host# set traps
```

To disable sending SNMP notifications on the physical interface, perform the following steps:

1. In configuration mode, go to the [edit interfaces *interface-name*] hierarchy level:

```
[edit]
user@host# edit interfaces interface-name
```

2. Configure the **no-traps** option to disable sending of Simple Network Management Protocol (SNMP) notifications when the state of the connection changes.

```
[edit interfaces interface-name]
user@host# set no-traps
```

Related Documentation

- [traps on page 262](#)

## Configuring Aggregated Ethernet Links (CLI Procedure)

Use the link aggregation feature to aggregate one or more links to form a virtual link or link aggregation group (LAG). The MAC client can treat this virtual link as if it were a single link to increase bandwidth, provide graceful degradation as failure occurs, and increase availability.



**NOTE:** An interface with an already configured IP address cannot form part of the aggregation group.

To configure aggregated Ethernet interfaces, using the CLI:

1. Specify the number of aggregated Ethernet interfaces to be created:

```
[edit chassis]
user@switch# set aggregated-devices ethernet device-count number
```

2. Specify the minimum number of links for the aggregated Ethernet interface (aex), that is, the defined bundle, to be labeled *up*:



**NOTE:** By default, only one link must be up for the bundle to be labeled *up*.

```
[edit interfaces]
user@switch# set ae0 aggregated-ether-options minimum-links number
```

3. Specify the link speed for the aggregated Ethernet bundle:

```
[edit interfaces]
user@switch# set ae0 aggregated-ether-options link-speed speed
```

4. Specify the members to be included within the aggregated Ethernet bundle:

```
[edit interfaces]
user@switch# set xe-fpc/pic/port ether-options 802.3ad ae0
user@switch# set xe-fpc/pic/port ether-options 802.3ad ae0
```

5. Specify an interface family for the aggregated Ethernet bundle:

```
[edit interfaces]
user@switch# set ae0 unit 0 family inet address address
```

For information about adding LACP to a LAG, see “Configuring Aggregated Ethernet LACP (CLI Procedure)” on page 113.

### Related Documentation

- [Configuring Aggregated Ethernet Interfaces \(J-Web Procedure\) on page 110](#)
- [Configuring Aggregated Ethernet LACP \(CLI Procedure\) on page 113](#)
- [Configuring LACP Link Protection of Aggregated Ethernet Interfaces \(CLI Procedure\) on page 114](#)
- [Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch on page 31](#)

- [Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch on page 37](#)
- [Verifying the Status of a LAG Interface on page 279](#)
- [Understanding Aggregated Ethernet Interfaces and LACP on page 8](#)

## Configuring Aggregated Ethernet Interfaces (J-Web Procedure)

---



**NOTE:** This topic applies only to the J-Web Application package.

Use the link aggregation feature to aggregate one or more Ethernet interfaces to form a virtual link or link aggregation group (LAG) on an EX Series switch. The MAC client can treat this virtual link as if it were a single link. Link aggregation increases bandwidth, provides graceful degradation as failure occurs, and increases availability. You can use the J-Web interface to configure aggregated Ethernet interfaces, or a LAG, on the switch.



**NOTE:** Interfaces that are already configured with MTU, duplex, flow control, or logical interfaces are listed but are not available for aggregation.

To configure an aggregated Ethernet interface (also referred to as a LAG):

1. Select **Configure > Interfaces > Link Aggregation**.

The list of aggregated interfaces is displayed.



**NOTE:** After you make changes to the configuration on this page, you must commit the changes immediately for them to take effect. To commit all changes to the active configuration, select **Commit Options > Commit**. See [Using the Commit Options to Commit Configuration Changes](#) for details about all commit options.

2. Click one of the following:

- **Add**—Creates an aggregated Ethernet interface, or LAG. Enter information as specified in [Table 17 on page 111](#).
- **Edit**—Modifies a selected LAG.
  - **Aggregation**—Modifies settings for the selected LAG. Enter information as specified in [Table 17 on page 111](#).
  - **VLAN**—Specifies VLAN options for the selected LAG. Enter information as specified in [Table 18 on page 112](#).
  - **IP Option**—Specifies IP options for the selected LAG. Enter information as specified in [Table 19 on page 112](#).

- **Delete**—Deletes the selected LAG.
- **Disable Port** or **Enable Port**—Disables or enables the administrative status on the selected interface.
- **Device Count**—Configures the number of aggregated logical devices available to the switch. Select the number and click **OK**.

Table 17: Aggregated Ethernet Interface Options

| Field                | Function                                                                                                                                                                                                                                                                                                                                                                                      | Your Action                                                                                                                                                                                                                                                                                                                                                                                                       |
|----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Aggregated Interface | Specifies the name of the aggregated interface.                                                                                                                                                                                                                                                                                                                                               | None. The name is supplied by the software.                                                                                                                                                                                                                                                                                                                                                                       |
| LACP Mode            | <p>Specifies the mode in which LACP packets are exchanged between the interfaces. The modes are:</p> <ul style="list-style-type: none"> <li>• <b>None</b>—Indicates that no mode is applicable.</li> <li>• <b>Active</b>—Indicates that the interface initiates transmission of LACP packets</li> <li>• <b>Passive</b>—Indicates that the interface responds only to LACP packets.</li> </ul> | Select from the list.                                                                                                                                                                                                                                                                                                                                                                                             |
| Description          | Specifies a description for the LAG.                                                                                                                                                                                                                                                                                                                                                          | Enter a description.                                                                                                                                                                                                                                                                                                                                                                                              |
| Interface            | Specifies the interfaces in the LAG.                                                                                                                                                                                                                                                                                                                                                          | <p>To add interfaces to the LAG, select the interfaces and click <b>Add</b>. For an EX8200 Virtual Chassis configuration, select the member, FPC, and the interface from the list. Click <b>OK</b>.</p> <p>To remove an interface from the LAG, select the interface and click <b>Remove</b>.</p> <p><b>NOTE:</b> Only interfaces that are configured with the same speed can be selected together for a LAG.</p> |
| Enable Log           | Specifies whether to enable generation of log entries for the LAG.                                                                                                                                                                                                                                                                                                                            | Select the check box to enable log generation, or clear the check box to disable log generation.                                                                                                                                                                                                                                                                                                                  |

Table 18: VLAN Options

| Field     | Function                                                       | Your Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|-----------|----------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Port Mode | Specifies the mode of operation for the port: trunk or access. | <p>If you select <b>Trunk</b>, you can:</p> <ol style="list-style-type: none"> <li>1. Click <b>Add</b> to add a VLAN member.</li> <li>2. Select the VLAN and click <b>OK</b>.</li> <li>3. (Optional) Associate a native VLAN ID with the port.</li> </ol> <p>If you select <b>Access</b>, you can:</p> <ol style="list-style-type: none"> <li>1. Select the VLAN member to be associated with the port.</li> <li>2. (Optional) Associate a VoIP VLAN with the interface. Only a VLAN with a VLAN ID can be associated as a VoIP VLAN.</li> </ol> <p>Click <b>OK</b>.</p> |

Table 19: IP Options

| Field        | Function                                        | Your Action                                                                                                                                                                                                                                                                                            |
|--------------|-------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| IPv4 Address | Specifies an IPv4 address for the selected LAG. | <ol style="list-style-type: none"> <li>1. Select the check box <b>IPv4 address</b>.</li> <li>2. Type an IP address—for example, <b>10.10.10.10</b>.</li> <li>3. Enter the subnet mask or address prefix. For example, 24 bits represents <b>255.255.255.0</b>.</li> <li>4. Click <b>OK</b>.</li> </ol> |
| IPv6 Address | Specifies an IPv6 address for the selected LAG. | <ol style="list-style-type: none"> <li>1. Select the check box <b>IPv6 address</b>.</li> <li>2. Type an IP address—for example, <b>2001:ab8:85a3::8a2e:370:7334</b>.</li> <li>3. Enter the subnet mask or address prefix.</li> <li>4. Click <b>OK</b>.</li> </ol>                                      |

**Related Documentation**

- [Configuring Aggregated Ethernet Links \(CLI Procedure\) on page 109](#)
- [Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch on page 31](#)
- [Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch on page 37](#)
- [Verifying the Status of a LAG Interface on page 279](#)

- [Configuring Aggregated Ethernet LACP \(CLI Procedure\)](#) on page 113
- [Understanding Aggregated Ethernet Interfaces and LACP](#) on page 8

## Configuring Aggregated Ethernet LACP (CLI Procedure)

For aggregated Ethernet interfaces on EX Series switches, you can configure the Link Aggregation Control Protocol (LACP). LACP is one method of bundling several physical interfaces to form one logical interface. You can configure aggregated Ethernet interfaces with or without LACP enabled.

LACP was designed to achieve the following:

- Automatic addition and deletion of individual links to the bundle without user intervention
- Link monitoring to check whether both ends of the bundle are connected to the correct group



**NOTE:** You can also configure LACP link protection on aggregated Ethernet interfaces. For information, see [“Configuring LACP Link Protection of Aggregated Ethernet Interfaces \(CLI Procedure\)”](#) on page 114.

The Junos OS implementation of LACP provides link monitoring but not automatic addition and deletion of links.

Before you configure LACP, be sure you have:

- Configured the aggregated Ethernet bundles—also known as link aggregation groups (LAGs). See [“Configuring Aggregated Ethernet Links \(CLI Procedure\)”](#) on page 109

When LACP is enabled, the local and remote sides of the aggregated Ethernet links exchange protocol data units (PDUs), which contain information about the state of the link. You can configure Ethernet links to actively transmit PDUs, or you can configure the links to passively transmit them (sending out LACP PDUs only when they receive them from another link). One side of the link must be configured as **active** for the link to be up.



**NOTE:** Do not add LACP to a LAG if the remote end of the LAG link is a security device, unless the security device supports LACP. Security devices often do not support LACP because they require a deterministic configuration.

To configure LACP:

1. Configure at least one side of the aggregated Ethernet link as active:

```
[edit interfaces]
user@switch# set aeX aggregated-ether-options lacp active
```

2. Specify the interval at which the interfaces send LACP packets:

```
[edit interfaces]
user@switch# set aeX aggregated-ether-options lacp periodic interval
```



**NOTE:** The LACP process exists in the system only if you configure the system in either active or passive LACP mode.

#### Related Documentation

- [Configuring Aggregated Ethernet Links \(CLI Procedure\) on page 109](#)
- [Configuring LACP Link Protection of Aggregated Ethernet Interfaces \(CLI Procedure\) on page 114](#)
- [Configuring Aggregated Ethernet Interfaces \(J-Web Procedure\) on page 110](#)
- [Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch on page 37](#)
- [Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch on page 31](#)
- [Verifying the Status of a LAG Interface on page 279](#)
- [Understanding Aggregated Ethernet Interfaces and LACP on page 8](#)

## Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure)

You can configure LACP link protection and system priority at the global level on the switch or for a specific aggregated Ethernet interface. When using LACP link protection to protect a single link in the aggregated ethernet bundle, you configure only two member links for an aggregated Ethernet interface: one active and one standby. LACP link protection ensures that only one link—the link with the higher priority—is used for traffic. The other link is forced to stay in a *waiting* state.

When using LACP link protection to protect multiple links in an aggregated ethernet bundle, you configure links into primary and backup subgroups. A link protection subgroup is a collection of ethernet links within the aggregated ethernet bundle. When you use link protection subgroups, you configure a primary subgroup and a backup subgroup. The configuration process includes assigning member links to each subgroup. When the configuration process is complete, the primary subgroup is used to forward traffic until a switchover event, such as a link failure, occurs and causes the backup subgroup to assume control of traffic that was travelling on the links in the primary subgroup within the bundle.

By default LACP link protection reverts to a higher-priority (lower-numbered) link when the higher-priority link becomes operational or when a higher-priority link is added to the aggregated Ethernet bundle. For priority purposes, LACP link protection treats subgroups like links. You can suppress link calculation by adding the **non-revertive** statement to the link protection configuration. In nonrevertive mode, when a link is active in sending and receiving LACP packets, adding a higher-priority link to the bundle does not change the status of the currently active link. It remains active.

If LACP link configuration is specified to be nonrevertive at the global **[edit chassis]** hierarchy level, you can specify the **revertive** statement in the LACP link protection configuration at the aggregated Ethernet interface level to override the nonrevertive setting for the interface. In revertive mode, adding a higher-priority link to the aggregated Ethernet bundle results in LACP recalculating the priority and switching the status from the currently active link to the newly added, higher-priority link.



**NOTE:** When LACP link protection is enabled on both local and remote sides of the link, both sides must use the same mode (either revertive or nonrevertive).

Configuring LACP link configuration at the aggregated Ethernet level results in only the configured interfaces using the defined configuration. LACP interface configuration also enables you to override global (chassis) LACP settings.

Before you configure LACP link protection, be sure you have:

- Configured the aggregated Ethernet bundles—also known as link aggregation groups (LAGs). See [“Configuring Aggregated Ethernet Links \(CLI Procedure\)” on page 109](#).
- Configured LACP for the interface. See [“Configuring Aggregated Ethernet LACP \(CLI Procedure\)” on page 113](#).

You can configure LACP link protection for all aggregated Ethernet interfaces on the switch by enabling it at the global level on the switch or configure it for a specific aggregated Ethernet interface by enabling it on that interface.

- [Configuring LACP Link Protection for a Single Link at the Global Level on page 116](#)
- [Configuring LACP Link Protection for a Single Link at the Aggregated Interface Level on page 116](#)
- [Configuring Subgroup Bundles to Provide LACP Link Protection to Multiple Links in an Aggregated Ethernet Interface on page 117](#)

## Configuring LACP Link Protection for a Single Link at the Global Level

To configure LACP link protection for aggregated Ethernet interfaces at the global level:

1. Enable LACP link protection on the switch:

```
[edit chassis aggregated-devices ethernet lacp]
user@switch# set link-protection
```

2. (Optional) Configure the LACP link protection for the aggregated Ethernet interfaces to be in nonrevertive mode:



**NOTE:** LACP link protection is in revertive mode by default.

```
[edit chassis aggregated-devices ethernet lacp link-protection]
user@switch# set non-revertive
```

3. (Optional) To configure LACP system priority for the aggregated Ethernet interfaces:

```
[edit chassis aggregated-devices ethernet lacp]
user@switch# set system-priority
```

## Configuring LACP Link Protection for a Single Link at the Aggregated Interface Level

To enable LACP link protection for a specific aggregated Ethernet interface:

1. Enable LACP link protection for the interface:

```
[edit interfaces aeX aggregated-ether-options lacp]
user@switch# set link-protection
```

2. (Optional) Configure the LACP link protection for the aggregated Ethernet interface to be in revertive or nonrevertive mode:

- To specify revertive mode:

```
[edit interfaces aeX aggregated-ether-options lacp link-protection]
user@switch# set revertive
```

- To specify nonrevertive mode:

```
[edit interfaces aeX aggregated-ether-options lacp link-protection]
user@switch# set non-revertive
```

3. (Optional) To configure LACP system priority for an aggregated Ethernet interface:

```
[edit interfaces aeX aggregated-ether-options lacp link-protection]
user@switch# set system-priority
```

4. (Optional) To configure LACP port priority for an aggregated Ethernet interface:

```
[edit interfaces ge-fpc/pic/port ether-options 802.3ad lacp]
user@switch# set port-priority
```

## Configuring Subgroup Bundles to Provide LACP Link Protection to Multiple Links in an Aggregated Ethernet Interface

You can configure link protection subgroup bundles to provide link protection for multiple links in an aggregated ethernet bundle.

Link protection subgroups allow you to provide link protection to a collection of Ethernet links within a LAG bundle, instead of providing protection to a single link in the aggregated ethernet bundle only. You can, for instance, configure a primary subgroup with three member links and a backup subgroup with three different member links and use the backup subgroup to provide link protection for the primary subgroup.

To configure link protection using subgroups:

1. Configure the primary link protection subgroup in the aggregated ethernet interface:

```
[edit interfaces aeX aggregated-ether-options]
user@switch# set link-protection-sub-group group-name primary
```

For instance, to create a primary link protection subgroup named **subgroup-primary** for interface **ae0**:

```
[edit interfaces ae0 aggregated-ether-options]
user@switch# set link-protection-sub-group subgroup-primary primary
```

2. Configure the backup link protection subgroup in the aggregated ethernet interface:

```
[edit interfaces aeX aggregated-ether-options]
user@switch# set link-protection-sub-group group-name backup
```

For instance, to create a backup link protection subgroup named **subgroup-backup** for interface **ae0**:

```
[edit interfaces ae0 aggregated-ether-options]
user@switch# set link-protection-sub-group subgroup-backup backup
```



**NOTE:** You can create one primary and one backup link protection subgroup per aggregated ethernet interface.

3. Attach interfaces to the link protection subgroups:

```
[edit interfaces interface-name ether-options 802.3ad]
user@switch# set link-protection-sub-group group-name
```



**NOTE:** The primary and backup link protection subgroups must contain the same number of interfaces. For instance, if the primary link protection subgroup contains three interfaces, the backup link protection subgroup must also contain three interfaces.

For instance, to configure interfaces **ge-0/0/0** and **ge-0/0/1** into link protection subgroup **subgroup-primary** and interfaces **ge-0/0/2** and **ge-0/0/3** into link protection subgroup **subgroup-backup**:

```
[edit interfaces ge-0/0/0 ether-options 802.3ad]
user@switch# set link-protection-sub-group subgroup-primary
```

- ```
[edit interfaces ge-0/0/1 ether-options 802.3ad]
user@switch# set link-protection-sub-group subgroup-primary
[edit interfaces ge-0/0/2 ether-options 802.3ad]
user@switch# set link-protection-sub-group subgroup-backup
[edit interfaces ge-0/0/3 ether-options 802.3ad]
user@switch# set link-protection-sub-group subgroup-backup
```
4. (Optional) Configure the port priority for link protection:

```
[edit interfaces interface-name ether-options 802.3ad]
user@switch# set port-priority priority
```

The port priority is used to select the active link.

5. Enable link protection

To enable link protection at the LAG level:

```
[edit interfaces aeX aggregated-ether-options]
user@switch# set link-protection
```

To enable link protection at the LACP level:

```
[edit interfaces aeX aggregated-ether-options lacp]
user@switch# set link-protection
```

For instance, to enable link protection on **ae0** at the LAG level:

```
[edit interfaces ae0 aggregated-ether-options]
user@switch# set link-protection
```

For instance, to enable link protection on **ae0** at the LACP level:

```
[edit interfaces ae0 aggregated-ether-options lacp]
user@switch# set link-protection
```

**Related
Documentation**

- [Understanding Aggregated Ethernet Interfaces and LACP on page 8](#)

Configuring Aggregated Ethernet Link Protection

You can configure link protection for aggregated Ethernet interfaces to provide QoS on the links during operation.

On aggregated Ethernet interfaces, you designate a primary and backup link to support link protection. Egress traffic passes only through the designated primary link. This includes transit traffic and locally generated traffic on the router or switch. When the primary link fails, traffic is routed through the backup link. Because some traffic loss is unavoidable, egress traffic is not automatically routed back to the primary link when the primary link is reestablished. Instead, you manually control when traffic should be diverted back to the primary link from the designated backup link.



NOTE: Link protection is not supported on MX80.

- [Configuring Link Protection for Aggregated Ethernet Interfaces on page 119](#)
- [Configuring Primary and Backup Links for Link Aggregated Ethernet Interfaces on page 119](#)

Configuring Aggregated Ethernet Link Speed

On aggregated Ethernet interfaces, you can set the required link speed for all interfaces included in the bundle. Generally, all interfaces that make up a bundle must have the same speed. If you include in the aggregated Ethernet interface an individual link that has a speed different from the speed that you specify in the **link-speed** parameter, an error message is logged. However, starting with Junos OS Release 13.2, aggregated Ethernet supports the following mixed rates and mixed modes on T640, T1600, T4000, and TX Matrix Plus routers:

- Member links of different modes (WAN and LAN) for 10-Gigabit Ethernet links.
- Member links of different rates: 10-Gigabit Ethernet, 40-Gigabit Ethernet, 50-Gigabit Ethernet, 100-Gigabit Ethernet, and OC192 (10-Gigabit Ethernet WAN mode)

Starting with Junos OS Release 14.2, aggregated Ethernet supports mixed link speeds on PTX Series Packet Transport Routers.



NOTE:

- Member links of 50-Gigabit Ethernet can only be configured using the 50-Gigabit Ethernet interfaces of 100-Gigabit Ethernet PIC with CFP (PD-1CE-CFP-FPC4).
- Starting with Junos OS Release 13.2, 100-Gigabit Ethernet member links can be configured using the two 50-Gigabit Ethernet interfaces of 100-Gigabit Ethernet PIC with CFP. This 100-Gigabit Ethernet member link can be included in an aggregated Ethernet link that includes member links of other interfaces as well. In releases before Junos OS Release 13.2, the 100-Gigabit Ethernet member link configured using the two 50-Gigabit Ethernet interfaces of 100-Gigabit Ethernet PIC with CFP cannot be included in an aggregated Ethernet link that includes member links of other interfaces.

To configure member links of mixed rates and mixed modes on T640, T1600, T4000, TX Matrix Plus, and PTX routers, you need to configure the **mixed** option for the **[edit interfaces aex aggregated-ether-options link-speed]** statement.

To set the required link speed:

1. Specify that you want to configure the aggregated Ethernet options.

```
user@host# edit interfaces interface-name aggregated-ether-options
```

2. Configure the link speed.

```
[edit interfaces interface-name aggregated-ether-options ]
user@host# set link-speed speed
```

speed can be in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000).

explicitly enabled, so even though unicast RPF is enabled on all interfaces, this status is not displayed.

- Enabling unicast RPF explicitly on all interfaces makes it easier to know whether unicast RPF is enabled on the switch because every interface shows the correct status. (Only interfaces on which you explicitly enable unicast RPF display the flag that indicates that unicast RPF is enabled.) The drawback of this approach is that if you want to disable unicast RPF, you must explicitly disable it on every interface. If unicast RPF is enabled on any interface, it is implicitly enabled on all interfaces.

Related Documentation

- [Example: Configuring Unicast RPF on an EX Series Switch on page 49](#)
- [Verifying Unicast RPF Status on page 281](#)
- [Disabling Unicast RPF \(CLI Procedure\) on page 125](#)
- [Troubleshooting Unicast RPF on page 429](#)
- [Understanding Unicast RPF on page 12](#)

Disabling Unicast RPF (CLI Procedure)

Unicast reverse-path forwarding (RPF) can help protect your LAN from denial-of-service (DoS) and distributed denial-of-service (DDoS) attacks on untrusted interfaces. Unicast RPF filters traffic with source addresses that do not use the incoming interface as the best return path back to the source. If the network configuration changes so that an interface that has unicast RPF enabled becomes a trusted interface or becomes asymmetrically routed (the interface that receives a packet is not the best return path to the packet's source), disable unicast RPF.

To disable unicast RPF on an EX3200, EX4200, or EX4300 switch, you must delete it from every interface on which you explicitly configured it. If you do not disable unicast RPF on every interface on which you explicitly enabled it, it remains implicitly enabled on all interfaces. If you attempt to delete unicast RPF from an interface on which it was not explicitly enabled, the **warning: statement not found** message appears. If you do not disable unicast RPF on every interface on which you explicitly enabled it, unicast RPF remains implicitly enabled on all interfaces of the EX3200, EX4200, or EX4300 switch.

On EX8200, EX6200, QFX Series switches, and OCX Series switches, the switch does not apply unicast RPF to an interface unless you explicitly enable that interface for unicast RPF.

CHAPTER 4

Configuration Statements

- [\[edit chassis\] Configuration Statement Hierarchy on EX Series Switches on page 139](#)
- [\[edit interfaces\] Configuration Statement Hierarchy on EX Series Switches on page 141](#)
- [\[edit interfaces ae\] Configuration Statement Hierarchy on EX Series Switches on page 142](#)
- [\[edit interfaces ge\] Configuration Statement Hierarchy on EX Series Switches on page 146](#)
- [\[edit interfaces interface-range\] Configuration Statement Hierarchy on EX Series Switches on page 149](#)
- [\[edit interfaces lo\] Configuration Statement Hierarchy on EX Series Switches on page 156](#)
- [\[edit interfaces me\] Configuration Statement Hierarchy on EX Series Switches on page 159](#)
- [\[edit interfaces vlan\] Configuration Statement Hierarchy on EX Series Switches on page 163](#)
- [\[edit interfaces vme\] Configuration Statement Hierarchy on EX Series Switches on page 166](#)
- [\[edit interfaces xe\] Configuration Statement Hierarchy on EX Series Switches on page 169](#)
- [\[edit protocols lacp\] Configuration Statement Hierarchy on EX Series Switches on page 173](#)
- [802.3ad on page 174](#)
- [accounting-profile on page 175](#)
- [address on page 176](#)
- [aggregated-devices on page 179](#)
- [aggregated-ether-options on page 180](#)
- [alarm \(optics-options\) on page 181](#)
- [arp \(Interfaces\) on page 182](#)
- [auto-negotiation on page 183](#)
- [bandwidth \(Interfaces\) on page 185](#)
- [broadcast on page 186](#)
- [chassis on page 187](#)

- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see [Feature Explorer](#).

This topic lists:

- [Supported Statements in the \[edit chassis\] Hierarchy Level on page 140](#)

Supported Statements in the [edit chassis] Hierarchy Level

The following hierarchy shows the **[edit chassis]** configuration statements supported on EX Series switches:

```
chassis {
  aggregated-devices {
    ethernet {
      device-count number;
      lacp {
        link-protection non-revertive;
        system-priority system-priority-number
      }
    }
  }
  alarm {
    ethernet {
      link-down (ignore | red | yellow);
    }
    management-ethernet {
      link-down (ignore | red | yellow);
    }
  }
  container-devices {
    device-count device-count-number;
  }
  disk-partition {
    /config {
      level (full | high) {
        free-space (free-space-threshold-value | mb | percent);
      }
    }
    /var {
      level (full | high) {
        free-space (free-space-threshold-value | mb | percent);
      }
    }
  }
  fpc slot-number {
    pic pic-number {
      no-multi-rate;
      q-pic-large-buffer (large-scale | small-scale);
    }
  }
}
```

```

maximum-ecmp maximum-ecmp-routes;
lcd-menu {
    fpc slot-number {
        menu-item menu-name);
        disable;
    }
    pseudowire-service {
        device-count device-count-number;
    }
    psu {
        redundancy {
            n-plus-n;
        }
        redundancy {
            graceful-switchover;
        }
    }
    slow-pfe-alarm;
}

```

**Related
Documentation**

- [Configuring Aggregated Ethernet Links \(CLI Procedure\) on page 109](#)
- [Configuring the LCD Panel on EX Series Switches \(CLI Procedure\)](#)
- [Configuring Graceful Routing Engine Switchover in a Virtual Chassis \(CLI Procedure\)](#)
- [Configuring Power Supply Redundancy \(CLI Procedure\)](#)
- [Configuring the Power Priority of Line Cards \(CLI Procedure\)](#)
- [Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade \(CLI Procedure\)](#)

[edit interfaces] Configuration Statement Hierarchy on EX Series Switches

Each of the following topics lists the statements at a subhierarchy of the **[edit interfaces]** hierarchy:

- [\[edit interfaces ae\] Configuration Statement Hierarchy on EX Series Switches on page 142](#)
- [\[edit interfaces ge\] Configuration Statement Hierarchy on EX Series Switches on page 146](#)
- [\[edit interfaces gr\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit interfaces interface-range\] Configuration Statement Hierarchy on EX Series Switches on page 149](#)
- [\[edit interfaces lo\] Configuration Statement Hierarchy on EX Series Switches on page 156](#)
- [\[edit interfaces me\] Configuration Statement Hierarchy on EX Series Switches on page 159](#)
- [\[edit interfaces vlan\] Configuration Statement Hierarchy on EX Series Switches on page 163](#)


```

lACP {
    (active | passive);
    admin-key key;
    periodic interval;
    system-id mac-address;
}
(link-protection | no-link-protection);
link-speed speed;
(loopback | no-loopback);
minimum-links number;
}
description text;
disable;
(gratuitous-arp-reply | no-gratuitous-arp-reply);
mtu bytes;
no-gratuitous-arp-request;
traceoptions {
    flag flag;
}
(traps | no-traps);
unit logical-unit-number {
    accounting-profile name;
    arp-resp;
    bandwidth rate;
    description text;
    disable;
    family ccc;
    family ethernet-switching {
        filter {
            input filter-name;
            output filter-name;
        }
        native-vlan-id vlan-id-number;
        port-mode (access | trunk);
        vlan {
            members [ members ];
        }
    }
}
family inet {
    address ipv4-address {
        arp ip-address (mac | multicast-mac) mac-address <publish>;
        broadcast address;
        preferred;
        primary;
        vrrp-group group-number {
            (accept-data | no-accept-data);
            advertise-interval seconds;
            authentication-key key;
            authentication-type authentication;
            fast-interval milliseconds;
            (preempt | no-preempt) {
                hold-time seconds;
            }
            priority number;
            track {
                interface interface-name {

```

```

        bandwidth-threshold bandwidth;
        priority-cost number;
    }
    priority-hold-time seconds;
    route ip-address/mask routing-instance instance-name priority-cost cost;
}
virtual-address [ addresses ];
virtual-link-local-address address;
vrrp-inherit-from {
    active-group group-number;
    active-interface interface-name;
}
}
}
dhcp {
    client-identifier (ascii client-id | hexadecimal client-id);
    lease-time (seconds | infinite);
    retransmission-attempt number;
    retransmission-interval seconds;
    server-address ip-address;
    update-server server;
    vendor-id id;
}
filter {
    input filter-name;
    output filter-name;
}
mtu bytes;
no-neighbor-learn;
no-redirects;
primary;
rpf-check;
targeted-broadcast;
}
family inet6 {
    address address {
        eui-64;
        ndp ip-address (mac | multicast-mac) mac-address <publish>;
        preferred;
        primary;
        vrrp-inet6-group group-id {
            accept-data | no-accept-data;
            authentication-key key;
            authentication-type authentication;
            fast-interval milliseconds;
            inet6-advertise-interval milliseconds;
            preempt | no-preempt {
                hold-time seconds;
            }
        }
        priority number;
        track {
            interface interface-name {
                bandwidth-threshold bandwidth priority-cost number;
                priority-cost number;
            }
        }
        priority-hold-time seconds;
    }
}

```

```

        route ( address | routing-instance routing-instance-name );
    }
    virtual-inet6-address [addresses];
    virtual-link-local-address ipv6-address;
}
    vrrp-inherit-from {
        active-group group-name;
        active-interface interface-name;
    }
}
(dad-disable | no-dad-disable);
filter {
    input filter-name;
    output filter-name;
}
mtu bytes;
no-neighbor-learn;
policer {
    input policer-name;
    output policer-name;
}
rpf-check;
}
family iso {
    address interface-address;
    mtu bytes;
}
family mpls {
    mtu bytes;
}
proxy-arp (restricted | unrestricted);
(traps | no-traps);
vlan-id (VLAN Tagging and Layer 3 Subinterfaces) vlan-id-number;
}
vlan-tagging;
}
}

```

Unsupported Statements in the [edit interfaces ae] Hierarchy Level

All statements in the [edit interfaces ae] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented with the following exceptions:

Table 20: Unsupported [edit interfaces ae] Configuration Statements on EX Series Switches

Statement	Hierarchy
NOTE: Variables, such as <i>interface-range</i> , are not shown in the statements or hierarchies.	
family fibre-channel	[edit interfaces ae unit]
source-address-filter	[edit interfaces ae aggregated-ether-options]
source-address-filtering no-source-address-filtering	[edit interfaces ae aggregated-ether-options]

- Related Documentation**
- [\[edit interfaces\] Configuration Statement Hierarchy on EX Series Switches on page 141](#)

[\[edit interfaces ge\] Configuration Statement Hierarchy on EX Series Switches](#)

This topic lists supported and unsupported configuration statements in the **[edit interfaces ge]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see [Feature Explorer](#).

This topic lists:

- [Supported Statements in the \[edit interfaces ge\] Hierarchy Level on page 146](#)
- [Unsupported Statements in the \[edit interfaces ge\] Hierarchy Level on page 149](#)

Supported Statements in the [edit interfaces ge] Hierarchy Level

The following hierarchy shows the **[edit interfaces ge]** configuration statements supported on EX Series switches.

```
interfaces {
  ge-fpc/pic/port {
    accounting-profile name;
    description text;
    disable;
    ether-options {
      802.3ad {
        aex;
        (backup | primary);
        lacp {
          force-up;
        }
      }
    }
    (auto-negotiation | no-auto-negotiation);
    (flow-control | no-flow-control);
    ieee-802-3az-eee;
    link-mode mode;
    (loopback | no-loopback);
    speed (auto-negotiation | speed);
  }
  (gratuitous-arp-reply | no-gratuitous-arp-reply);
  hold-time up milliseconds down milliseconds;
  mtu bytes;
  no-gratuitous-arp-request;
  optics-options {
    alarm alarm-type;
    warning alarm-type;
```

```

    wavelength nanometers;
}
traceoptions {
    flag flag;
}
(traps | no-traps);
unit logical-unit-number {
    accounting-profile name;
    arp-resp;
    bandwidth rate;
    description text;
    disable;
    family ccc;
    family ethernet-switching {
        filter {
            input filter-name;
            output filter-name;
        }
        native-vlan-id vlan-id-number;
        port-mode (access | trunk);
        vlan {
            members [ members ];
        }
    }
}
family inet {
    address ipv4-address {
        arp ip-address (mac | multicast-mac) mac-address <publish>;
        broadcast address;
        preferred;
        primary;
        vrrp-group group-number {
            (accept-data | no-accept-data);
            advertise-interval seconds;
            authentication-key key;
            authentication-type authentication;
            fast-interval milliseconds;
            (preempt | no-preempt) {
                hold-time seconds;
            }
        }
        priority number;
        track {
            interface interface-name {
                bandwidth-threshold bandwidth;
                priority-cost number;
            }
            priority-hold-time seconds;
            route ip-address/mask routing-instance instance-name priority-cost cost;
        }
        virtual-address [ addresses ];
        virtual-link-local-address address;
        vrrp-inherit-from {
            active-group group-number;
            active-interface interface-name;
        }
    }
}
}

```

```

dhcp {
  client-identifier (ascii client-id | hexadecimal client-id);
  lease-time (seconds | infinite);
  retransmission-attempt number;
  retransmission-interval seconds;
  server-address ip-address;
  update-server
  vendor-id
}
filter {
  input filter-name;
  output filter-name;
}
mtu bytes;
no-neighbor-learn;
no-redirects;
primary;
rpf-check;
targeted-broadcast;
}
family inet6 {
  address address {
    eui-64;
    ndp ip-address (mac | multicast-mac) mac-address <publish>;
    preferred;
    primary;
    vrrp-inet6-group group-id {
      accept-data | no-accept-data;
      authentication-key key;
      authentication-type authentication;
      fast-interval milliseconds;
      inet6-advertise-interval milliseconds;
      preempt | no-preempt {
        hold-time seconds;
      }
      priority number;
      track {
        interface interface-name {
          bandwidth-threshold bandwidth priority-cost number;
          priority-cost number;
        }
        priority-hold-time seconds;
        route ( address | routing-instance routing-instance-name );
      }
      virtual-inet6-address [addresses];
      virtual-link-local-address ipv6-address;
      vrrp-inherit-from {
        active-group group-name;
        active-interface interface-name;
      }
    }
  }
}
(dad-disable | no-dad-disable);
filter {
  group group-name;
  input filter-name;
}

```

```

        output filter-name;
    }
    mtu bytes;
    no-neighbor-learn;
    policer {
        input policer-name;
        output policer-name;
    }
    rpf-check;
}
family iso {
    address interface-address;
    mtu bytes;
}
family mpls {
    mtu bytes;
}
proxy-arp (restricted | unrestricted);
swap-by-poppush;
(traps | no-traps);
vlan-id vlan-id-number;
}
vlan-tagging;
}
}

```

Unsupported Statements in the [edit interfaces ge] Hierarchy Level

All statements in the [edit interfaces ge] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented.

Related Documentation • [\[edit interfaces\] Configuration Statement Hierarchy on EX Series Switches on page 141](#)

[edit interfaces interface-range] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the [edit interfaces interface-range] hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see [Feature Explorer](#).

This topic lists:

- [Supported Statements in the \[edit interfaces interface-range\] Hierarchy Level on page 150](#)
- [Unsupported Statements in the \[edit interfaces interface-range\] Hierarchy Level on page 154](#)

Supported Statements in the [edit interfaces interface-range] Hierarchy Level

The following hierarchy shows the **[edit interfaces interface-range]** configuration statements supported on EX Series switches.

```
interfaces {
  interface-range name {
    accounting-profile name;
    aggregated-ether-options {
      ethernet-switch-profile {
        tag-protocol-id identifier;
      }
      (flow-control | no-flow-control);
      ieee-802-3ad eee;
      lacp {
        (active | passive);
        admin-key key;
        periodic interval;
        system-id mac-address;
      }
      (link-protection | no-link-protection);
      link-speed speed;
      (loopback | no-loopback);
      minimum-links number;
      rebalance-periodic;
      source-address-filter filter;
      source-filtering | no-source-filtering;
    }
    description text;
    disable;
    ether-options {
      802.3ad {
        aex;
        (backup | primary);
        lacp {
          force-up;
        }
      }
    }
    (auto-negotiation | no-auto-negotiation);
    (flow-control | no-flow-control);
    link-mode mode;
    (loopback | no-loopback);
    speed (auto-negotiation | speed);
  }
  framing;
  (gratuitous-arp-reply | no-gratuitous-arp-reply);
  hold-time up milliseconds down milliseconds;
  member interface-name;
```

```

member-range starting-interface name to ending-interface name;
mtu bytes;
no-gratuitous-arp-request;
optics-options {
    alarm alarm-type;
    warning alarm-type;
    wavelength nanometers;
}
services-options;
speed speed;
traceoptions {
    flag flag;
}
(traps | no-traps);
unit logical-unit-number {
    accept-source-mac;
    accounting-profile name;
    arp-resp;
    bandwidth rate;
    description text;
    disable;
    family ccc;
    family ethernet-switching {
        filter {
            input filter-name;
            output filter-name;
        }
        native-vlan-id vlan-id-number;
        port-mode (access | trunk);
        vlan {
            members [ members ];
        }
    }
}
family inet {
    accounting {
        destination-class-usage;
        source-class-usage;
    }
    address ipv4-address {
        arp ip-address (mac | multicast-mac) mac-address <publish>;
        broadcast address;
        destination-class-usage;
        destination-profile;
        master-only;
        preferred;
        primary;
        vrrp-group group-number {
            (accept-data | no-accept-data);
            advertise-interval seconds;
            authentication-key key;
            authentication-type authentication;
            fast-interval milliseconds;
            (preempt | no-preempt) {
                hold-time seconds;
            }
            priority number;
        }
    }
}

```

```

track {
    interface interface-name {
        bandwidth-threshold bandwidth;
        priority-cost number;
    }
    priority-hold-time seconds;
    route ip-address/mask routing-instance instance-name priority-cost cost;
}
virtual-address [ addresses ];
virtual-link-local-address address;
vrrp-inherit-from {
    active-group group-number;
    active-interface interface-name;
}
}
}
dhcp {
    client-identifier (ascii client-id | hexadecimal client-id);
    lease-time (seconds | infinte);
    retransmission-attempt number;
    retransmission-interval sections;
    server-address ip-address;
    update-server
    vendor-id
}
filter {
    input filter-name;
    output filter-name;
}
ipsec-sa;
mtu bytes;
multicast-only;
negotiate-address;
next-hop-tunnel;
no-neighbor-learn;
no-redirects;
primary;
receive-option-packets;
rpf-check;
targeted-broadcast;
}
family inet6 {
    accounting {
        destination-class-usage;
        source-class-usage;
    }
    address address {
        eui-64;
        ndp ip-address (mac | multicast-mac) mac-address <publish>;
        preferred;
        primary;
        vrrp-inet6-group group-id {
            accept-data | no-accept-data;
            authentication-key key;
            authentication-type authentication;
            fast-interval milliseconds;

```

```

    inet6-advertise-interval milliseconds;
    preempt | no-preempt {
        hold-time seconds;
    }
    priority number;
    track {
        interface interface-name {
            bandwidth-threshold bandwidth priority-cost number;
            priority-cost number;
        }
        priority-hold-time seconds;
        route ( address | routing-instance routing-instance-name );
    }
    virtual-inet6-address [addresses];
    virtual-link-local-address ipv6-address;
}
vrrp-inherit-from {
    active-group group-name;
    active-interface interface-name;
}
}
(dad-disable | no-dad-disable);
filter {
    group group-name;
    input filter-name;
    output filter-name;
}
mtu bytes;
no-neighbor-learn;
policer {
    input policer-name;
    output policer-name;
}
rpf-check;
}
family iso {
    address interface-address;
    mtu bytes;
}
family mpls {
    mtu bytes;
}
interface-shared-with;
interleave-fragments;
inverse-arp;
link-layer-overhead;
minimum-links;
mtu;
proxy-arp (restricted | unrestricted);
swap-by-poppush;
(traps | no-traps);
vlan-id vlan-id-number;
}
vlan-tagging;
}

```

Unsupported Statements in the [edit interfaces interface-range] Hierarchy Level

All statements in the [edit interfaces interface-range] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented with the following exceptions:

Table 21: Unsupported [edit interfaces interface-range] Configuration Statements for EX Series Switches

Statement	Hierarchy
NOTE: Variables, such as <i>interface-range</i> , are not shown in the statements or hierarchies.	
aggregated-sonet-options and all substatements	[edit interfaces interface-range]
allow-any-vci	[edit interfaces interface-range unit]
atm-l2circuit-mode	[edit interfaces interface-range unit]
atm-options and all substatements	[edit interfaces interface-range]
cell-bundle-size	[edit interfaces interface-range unit]
clear-don't-fragment-bit	[edit interfaces interface-range unit]
clocking	[edit interfaces interface-range]
compression-device	[edit interfaces interface-range unit]
container-options and all substatements	[edit interfaces interface-range]
copy-tos-to-outer-ip-header	[edit interfaces interface-range unit]
dce	[edit interfaces interface-range]
disable-mlppp-inter-ppp-pfc	[edit interfaces interface-range unit]
dlci	[edit interfaces interface-range unit]
drop-timeout	[edit interfaces interface-range unit]
ds0-options and all substatements	[edit interfaces interface-range]
e1-options and all substatements	[edit interfaces interface-range]
e3-options and all substatements	[edit interfaces interface-range]
epd-threshold	[edit interfaces interface-range unit]
family mlfr-end-to-end and all substatements	[edit interfaces interface-range unit]

Table 21: Unsupported [edit interfaces interface-range] Configuration Statements for EX Series Switches (*continued*)

Statement	Hierarchy
family mlfr-uni-uni and all substatements	[edit interfaces interface-range unit]
family mlppp and all substatements	[edit interfaces interface-range unit]
fragment-threshold	[edit interfaces interface-range unit]
ggsn-options and all substatements	[edit interfaces interface-range]
keepalives no-keepalives	[edit interfaces interface-range] [edit interfaces interface-range unit]
lmi	[edit interfaces interface-range]
lsq-failure-options	[edit interfaces interface-range]
mlfr-uni-nni-bundle-options and all substatements	[edit interfaces interface-range]
mrru	[edit interfaces interface-range unit]
multicast-dlci	[edit interfaces interface-range unit]
multilink-max-classes	[edit interfaces interface-range unit]
multipoint	[edit interfaces interface-range unit]
multipoint-destination	[edit interfaces interface-range unit family inet address]
multiservice-options and all substatements	[edit interfaces interface-range]
oam-liveness	[edit interfaces interface-range unit]
oam-period	[edit interfaces interface-range unit]
passive-monitor-mode	[edit interfaces interface-range unit]
peer-unit	[edit interfaces interface-range unit]
plp-to-clp	[edit interfaces interface-range unit]
point-to-point	[edit interfaces interface-range unit]
ppp-options and all substatements	[edit interfaces interface-range] [edit interfaces interface-range unit]
receive-lsp	[edit interfaces interface-range unit]

Table 21: Unsupported [edit interfaces interface-range] Configuration Statements for EX Series Switches (*continued*)

Statement	Hierarchy
satop-options and all substatements	[edit interfaces interface-range]
serial-options and all substatements	[edit interfaces interface-range]
service-domain	[edit interfaces interface-range unit]
shaping	[edit interfaces interface-range unit]
short-sequence	[edit interfaces interface-range unit]
shdsl-options and all substatements	[edit interfaces interface-range]
t1-options and all substatements	[edit interfaces interface-range]
t3-options and all substatements	[edit interfaces interface-range]
transmit-lsp	[edit interfaces interface-range unit]
transmit-weight	[edit interfaces interface-range unit]
trunk-id	[edit interfaces interface-range unit]
tunnel	[edit interfaces interface-range unit]
vci	[edit interfaces interface-range unit]
vci-range	[edit interfaces interface-range unit]
vpi	[edit interfaces interface-range unit]

Related Documentation • [\[edit interfaces\] Configuration Statement Hierarchy on EX Series Switches on page 141](#)

[\[edit interfaces lo\] Configuration Statement Hierarchy on EX Series Switches](#)

This topic lists supported and unsupported configuration statements in the **[edit interfaces lo]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see [Feature Explorer](#).

This topic lists:

- [Supported Statements in the \[edit interfaces lo\] Hierarchy Level on page 157](#)
- [Unsupported Statements in the \[edit interfaces lo\] Hierarchy Level on page 159](#)

Supported Statements in the [edit interfaces lo] Hierarchy Level

The following hierarchy shows the [edit interfaces lo] configuration statements supported on EX Series switches.

```

interfaces {
  lo0 {
    accounting-profile name;
    description text;
    disable;
    hold-time down milliseconds up milliseconds ;
    traceoptions {
      flag flag;
    }
    (traps | no-traps);
    unit logical-unit-number {
      accounting-profile name;
      arp-resp;
      bandwidth rate;
      description text;
      disable;
      family ccc;
      family inet {
        address ipv4-address {
          preferred;
          primary;
        }
        vrrp-group group-number {
          (accept-data | no-accept-data);
          advertise-interval seconds;
          authentication-key key;
          authentication-type authentication;
          fast-interval milliseconds;
          (preempt | no-preempt) {
            hold-time seconds;
          }
          priority number;
        }
        track {
          interface interface-name {
            bandwidth-threshold bandwidth;
            priority-cost number;
          }
          priority-hold-time seconds;
          route ip-address/mask routing-instance instance-name priority-cost cost;
        }
        virtual-address [ addresses ];
        virtual-link-local-address address;
        vrrp-inherit-from {
          active-group group-number;
          active-interface interface-name;
        }
      }
    }
  }
}

```

```

    }
}
dhcp {
    client-identifier (ascii client-id | hexadecimal client-id);
    lease-time (seconds | infinite);
    retransmission-attempt number;
    retransmission-interval seconds;
    server-address ip-address;
    update-server
    vendor-id
}
filter {
    input filter-name;
    output filter-name;
}
no-neighbor-learn;
no-redirects;
primary;
}
family inet6 {
    address address {
        preferred;
        primary;
        vrrp-inet6-group group-id {
            accept-data | no-accept-data;
            authentication-key key;
            authentication-type authentication;
            fast-interval milliseconds;
            inet6-advertise-interval milliseconds;
            preempt | no-preempt {
                hold-time seconds;
            }
            priority number;
            track {
                interface interface-name {
                    bandwidth-threshold bandwidth priority-cost number;
                    priority-cost number;
                }
                priority-hold-time seconds;
                route ( address | routing-instance routing-instance-name );
            }
            virtual-inet6-address [addresses];
            virtual-link-local-address ipv6-address;
            vrrp-inherit-from {
                active-group group-name;
                active-interface interface-name;
            }
        }
    }
    (dad-disable | no-dad-disable);
    filter {
        group group-name;
        input filter-name;
        output filter-name;
    }
    no-neighbor-learn;
    policer {

```

```

        input policer-name;
        output policer-name;
    }
}
family iso {
    address interface-address;
}
family mpls;
(traps | no-traps);
}
}
}

```

Unsupported Statements in the [edit interfaces lo] Hierarchy Level

All statements in the [edit interfaces lo] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented with the following exceptions:

Table 22: Unsupported [edit interfaces lo] Configuration Statements for EX Series Switches

Statement	Hierarchy
layer2-policer	[edit interfaces lo unit]
any	[edit interfaces lo unit family]
tcc	[edit interfaces lo unit family]
policer	[edit interfaces lo unit family inet]
unnumbered-address	[edit interfaces lo unit family inet]

- Related Documentation**
- [\[edit interfaces\] Configuration Statement Hierarchy on EX Series Switches on page 141](#)
 - [\[edit interfaces\] Configuration Statement Hierarchy on EX Series Switches](#)

[edit interfaces me] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the [edit interfaces me] hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see [Feature Explorer](#).

This topic lists:

- [Supported Statements in the \[edit interfaces me\] Hierarchy Level on page 160](#)
- [Unsupported Statements in the \[edit interfaces me\] Hierarchy Level on page 162](#)

Supported Statements in the [edit interfaces me] Hierarchy Level

The following hierarchy shows the **[edit interfaces me]** configuration statements supported on EX Series switches.

```
interfaces {
  me0 {
    accounting-profile name;
    description text;
    disable;
    (gratuitous-arp-reply | no-gratuitous-arp-reply);
    hold-time up milliseconds down milliseconds;
    no-gratuitous-arp-request;
    traceoptions {
      flag flag;
    }
    (traps | no-traps);
    unit logical-unit-number {
      accounting-profile name;
      arp-resp;
      bandwidth rate;
      description text;
      disable;
      family ethernet-switching {
        filter {
          input filter-name;
          output filter-name;
        }
        native-vlan-id vlan-id-number;
        port-mode (access | trunk);
        vlan {
          members [ members ];
        }
      }
    }
    family inet {
      accounting {
        destination-class-usage;
        source-class-usage {
          input;
          output;
        }
      }
      address ipv4-address {
        arp ip-address (mac | multicast-mac) mac-address <publish>;
        broadcast address;
        master-only;
        preferred;
        primary;
      }
      dhcp {
```

```

    client-identifier (ascii client-id | hexadecimal client-id);
    lease-time (seconds | infinte);
    retransmission-attempt number;
    retransmission-interval sections;
    server-address ip-address;
    update-server
    vendor-id
}
filter {
    input filter-name;
    output filter-name;
}
mtu bytes;
no-neighbor-learn;
primary;
rpf-check;
}
family inet6 {
    accounting {
        destination-class-usage;
        source-class-usage {
            input;
            output;
        }
    }
}
address address {
    eui-64;
    ndp ip-address (mac | multicast-mac) mac-address <publish>;
    preferred;
    primary;
}
(dad-disable | no-dad-disable);
filter {
    group group-name;
    input filter-name;
    output filter-name;
}
mtu bytes;
no-neighbor-learn;
policer {
    input policer-name;
    output policer-name;
}
rpf-check;
}
family iso {
    address interface-address;
    mtu bytes;
}
family mpls {
    mtu bytes;
}
swap-by-poppush;
(traps | no-traps);
vlan-id vlan-id-number;
}

```

```

        vlan-tagging;
    }
}

```

Unsupported Statements in the [edit interfaces me] Hierarchy Level

All statements in the **[edit interfaces me]** hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented with the following exceptions:

Table 23: Unsupported [edit interfaces me] Configuration Statements for EX Series Switches

Statement	Hierarchy
encapsulation	[edit interfaces me]
link-mode	[edit interfaces me]
encapsulation	[edit interfaces me unit]
layer2-policer	[edit interfaces me unit]
native-inner-vlan-id	[edit interfaces me unit]
vlan-id-list	[edit interfaces me unit]
vlan-id-range	[edit interfaces me unit]
ccc	[edit interfaces me unit family]
tcc	[edit interfaces me unit family]
vpls	[edit interfaces me unit family]
no-redirects	[edit interfaces me unit family inet]
policer	[edit interfaces me unit family inet]
sampling	[edit interfaces me unit family inet]
service	[edit interfaces me unit family inet]
unnumbered-address	[edit interfaces me unit family inet]
vrrp-group	[edit interfaces me unit family inet address]
service	[edit interfaces me unit family inet6]
vrrp-inet6-group	[edit interfaces me unit family inet6 address]

- Related Documentation**
- [\[edit interfaces\] Configuration Statement Hierarchy on EX Series Switches on page 141](#)
 - [\[edit interfaces\] Configuration Statement Hierarchy on EX Series Switches](#)

[\[edit interfaces vlan\] Configuration Statement Hierarchy on EX Series Switches](#)

This topic lists supported and unsupported configuration statements in the **[edit interfaces vlan]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see [Feature Explorer](#).

This topic lists:

- [Supported Statements in the \[edit interfaces vlan\] Hierarchy Level on page 163](#)
- [Unsupported Statements in the \[edit interfaces vlan\] Hierarchy Level on page 166](#)

Supported Statements in the **[edit interfaces vlan]** Hierarchy Level

The following hierarchy shows the **[edit interfaces vlan]** configuration statements supported on EX Series switches.

```

interfaces {
  vlan {
    accounting-profile name;
    description text;
    disable;
    (gratuitous-arp-reply | no-gratuitous-arp-reply);
    hold-time up milliseconds down milliseconds;
    mtu bytes;
    no-gratuitous-arp-request;
    traceoptions {
      flag flag;
    }
    (traps | no-traps);
    unit logical-unit-number {
      accounting-profile name;
      arp-resp;
      bandwidth rate;
      description text;
      disable;
      family inet {
        accounting {
          destination-class-usage;
          source-class-usage {
            input;
            output;
          }
        }
      }
    }
  }
}

```

```

}
address ipv4-address {
  arp ip-address (mac | multicast-mac) mac-address <publish>;
  broadcast address;
  master-only;
  preferred;
  primary;
  vrrp-group group-number {
    (accept-data | no-accept-data);
    advertise-interval seconds;
    authentication-key key;
    authentication-type authentication;
    fast-interval milliseconds;
    (preempt | no-preempt) {
      hold-time seconds;
    }
    priority number;
    track {
      interface interface-name {
        bandwidth-threshold bandwidth;
        priority-cost number;
      }
      priority-hold-time seconds;
      route ip-address/mask routing-instance instance-name priority-cost cost;
    }
    virtual-address [ addresses ];
    virtual-link-local-address address;
    vrrp-inherit-from {
      active-group group-number;
      active-interface interface-name;
    }
  }
}
}
dhcp {
  client-identifier (ascii client-id | hexadecimal client-id);
  lease-time (seconds | infinite);
  retransmission-attempt number;
  retransmission-interval sections;
  server-address ip-address;
  update-server
  vendor-id
}
filter {
  input filter-name;
  output filter-name;
}
mtu bytes;
no-neighbor-learn;
primary;
rpf-check;
}
family inet6 {
  accounting {
    destination-class-usage;
    source-class-usage {
      input;
    }
  }
}

```

```

        output;
    }
}
address address {
    eui-64;
    ndp ip-address (mac | multicast-mac) mac-address <publish>;
    preferred;
    primary;
    vrrp-inet6-group group-id {
        accept-data | no-accept-data;
        authentication-key key;
        authentication-type authentication;
        fast-interval milliseconds;
        inet6-advertise-interval milliseconds;
        preempt | no-preempt {
            hold-time seconds;
        }
        priority number;
        track {
            interface interface-name {
                bandwidth-threshold bandwidth priority-cost number;
                priority-cost number;
            }
            priority-hold-time seconds;
            route ( address | routing-instance routing-instance-name );
        }
        virtual-inet6-address [addresses];
        virtual-link-local-address ipv6-address;
        vrrp-inherit-from {
            active-group group-name;
            active-interface interface-name;
        }
    }
}
(dad-disable | no-dad-disable);
filter {
    group group-name;
    input filter-name;
    output filter-name;
}
mtu bytes;
no-neighbor-learn;
policer {
    input policer-name;
    output policer-name;
}
rpf-check;
}
family iso {
    address interface-address;
    mtu bytes;
}
family mpls {
    mtu bytes;
}
proxy-arp (restricted | unrestricted);

```

```
        (traps | no-traps);
    }
}
```

Unsupported Statements in the [edit interfaces vlan] Hierarchy Level

All statements in the **[edit interfaces vlan]** hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented.

Related Documentation

- [\[edit interfaces\] Configuration Statement Hierarchy on EX Series Switches on page 141](#)

[edit interfaces vme] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the **[edit interfaces vme]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see [Feature Explorer](#).

This topic lists:

- [Supported Statements in the \[edit interfaces vme\] Hierarchy Level on page 166](#)
- [Unsupported Statements in the \[edit interfaces vme\] Hierarchy Level on page 169](#)

Supported Statements in the [edit interfaces vme] Hierarchy Level

The following hierarchy shows the **[edit interfaces vme]** configuration statements supported on EX Series switches.

```
interfaces {
  vme {
    accounting-profile name;
    description text;
    disable;
    (gratuitous-arp-reply | no-gratuitous-arp-reply);
    hold-time up milliseconds down milliseconds;
    mtu bytes;
    no-gratuitous-arp-request;
    traceoptions {
      flag flag;
    }
    (traps | no-traps);
    unit logical-unit-number {
      accounting-profile name;
      arp-resp;
      bandwidth rate;
```

```

description text;
disable;
family inet {
    accounting {
        destination-class-usage;
        source-class-usage {
            input;
            output;
        }
    }
}
address ipv4-address {
    arp ip-address (mac | multicast-mac) mac-address <publish>;
    broadcast address;
    master-only;
    preferred;
    primary;
    vrrp-group group-number {
        (accept-data | no-accept-data);
        advertise-interval seconds;
        authentication-key key;
        authentication-type authentication;
        fast-interval milliseconds;
        (preempt | no-preempt) {
            hold-time seconds;
        }
        priority number;
        track {
            interface interface-name {
                bandwidth-threshold bandwidth;
                priority-cost number;
            }
            priority-hold-time seconds;
            route ip-address/mask routing-instance instance-name priority-cost cost;
        }
        virtual-address [ addresses ];
        virtual-link-local-address address;
        vrrp-inherit-from {
            active-group group-number;
            active-interface interface-name;
        }
    }
}
dhcp {
    client-identifier (ascii client-id | hexadecimal client-id);
    lease-time (seconds | infinite);
    retransmission-attempt number;
    retransmission-interval seconds;
    server-address ip-address;
    update-server
    vendor-id
}
filter {
    input filter-name;
    output filter-name;
}
mtu bytes;

```

```

no-neighbor-learn;
primary;
rpf-check;
}
family inet6 {
  accounting {
    destination-class-usage;
    source-class-usage {
      input;
      output;
    }
  }
}
address address {
  eui-64;
  ndp ip-address (mac | multicast-mac) mac-address <publish>;
  preferred;
  primary;
  vrrp-inet6-group group-id {
    accept-data | no-accept-data;
    authentication-key key;
    authentication-type authentication;
    fast-interval milliseconds;
    inet6-advertise-interval milliseconds;
    preempt | no-preempt {
      hold-time seconds;
    }
    priority number;
    track {
      interface interface-name {
        bandwidth-threshold bandwidth priority-cost number;
        priority-cost number;
      }
      priority-hold-time seconds;
      route ( address | routing-instance routing-instance-name );
    }
    virtual-inet6-address [addresses];
    virtual-link-local-address ipv6-address;
    vrrp-inherit-from {
      active-group group-name;
      active-interface interface-name;
    }
  }
}
(dad-disable | no-dad-disable);
filter {
  group group-name;
  input filter-name;
  output filter-name;
}
mtu bytes;
no-neighbor-learn;
policer {
  input policer-name;
  output policer-name;
}
rpf-check;

```

```

    }
    family iso {
        address interface-address;
        mtu bytes;
    }
    family mpls {
        mtu bytes;
    }
    (traps | no-traps);
    vlan-id vlan-id-number;
}
vlan-tagging;
}
}

```

Unsupported Statements in the [edit interfaces vme] Hierarchy Level

All statements in the [edit interfaces vme] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented.

- Related Documentation**
- [\[edit interfaces\] Configuration Statement Hierarchy on EX Series Switches on page 141](#)
 - [\[edit interfaces\] Configuration Statement Hierarchy on EX Series Switches](#)

[edit interfaces xe] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the [edit interfaces xe] hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see [Feature Explorer](#).

This topic lists:

- [Supported Statements in the \[edit interfaces xe\] Hierarchy Level on page 169](#)
- [Unsupported Statements in the \[edit interfaces xe\] Hierarchy Level on page 173](#)

Supported Statements in the [edit interfaces xe] Hierarchy Level

The following hierarchy shows the [edit interfaces xe] configuration statements supported on EX Series switches.

```

interfaces {
  xe-fpc/pic/port {
    accounting-profile name;
    clocking (external | internal);
    description text;
  }
}

```

```

disable;
ether-options {
    802.3ad {
        aex;
        (backup | primary);
        lacp {
            force-up;
        }
    }
    (flow-control | no-flow-control);
    (loopback | no-loopback);
}
framing (lan-phy | wan-phy);
(gratuitous-arp-reply | no-gratuitous-arp-reply);
hold-time up milliseconds down milliseconds;
mtu bytes;
no-gratuitous-arp-request;
optics-options {
    alarm alarm-type;
    warning alarm-type;
    wavelength nanometers;
}
traceoptions {
    flag flag;
}
(traps | no-traps);
unit logical-unit-number {
    accounting-profile name;
    bandwidth rate;
    description text;
    disable;
    family ccc;
    family ethernet-switching {
        reflective-relay
        filter {
            input filter-name;
            output filter-name;
        }
        native-vlan-id vlan-id-number;
        port-mode (access | trunk | tagged-access);
        vlan {
            members [ members ];
        }
    }
}
family inet {
    accounting {
        destination-class-usage;
        source-class-usage {
            input;
            output;
        }
    }
}
address ipv4-address {
    arp ip-address (mac | multicast-mac) mac-address <publish>;
    broadcast address;
    preferred;

```

```

primary;
vrrp-group group-number {
  (accept-data | no-accept-data);
  advertise-interval seconds;
  authentication-key key;
  authentication-type authentication;
  fast-interval milliseconds;
  (preempt | no-preempt) {
    hold-time seconds;
  }
  priority number;
  track {
    interface interface-name {
      bandwidth-threshold bandwidth;
      priority-cost number;
    }
    priority-hold-time seconds;
    route ip-address/mask routing-instance instance-name priority-cost cost;
  }
  virtual-address [ addresses ];
  virtual-link-local-address address;
  vrrp-inherit-from {
    active-group group-number;
    active-interface interface-name;
  }
}
}
dhcp {
  client-identifier (ascii client-id | hexadecimal client-id);
  lease-time (seconds | infinte);
  retransmission-attempt number;
  retransmission-interval sections;
  server-address ip-address;
  update-server
  vendor-id
}
filter {
  input filter-name;
  output filter-name;
}
mtu bytes;
no-neighbor-learn;
no-redirects;
primary;
rpf-check;
targeted-broadcast;
}
family inet6 {
  accounting {
    destination-class-usage;
    source-class-usage {
      input;
      output;
    }
  }
}
address address {

```

```

eui-64;
ndp ip-address (mac | multicast-mac) mac-address <publish>;
preferred;
primary;
vrrp-inet6-group group-id {
    accept-data | no-accept-data;
    authentication-key key;
    authentication-type authentication;
    fast-interval milliseconds;
    inet6-advertise-interval milliseconds;
    preempt | no-preempt {
        hold-time seconds;
    }
    priority number;
    track {
        interface interface-name {
            bandwidth-threshold bandwidth priority-cost number;
            priority-cost number;
        }
        priority-hold-time seconds;
        route ( address | routing-instance routing-instance-name );
    }
    virtual-inet6-address [addresses];
    virtual-link-local-address ipv6-address;
    vrrp-inherit-from {
        active-group group-name;
        active-interface interface-name;
    }
}
}
(dad-disable | no-dad-disable);
filter {
    group group-name;
    input filter-name;
    output filter-name;
}
mtu bytes;
no-neighbor-learn;
policer {
    input policer-name;
    output policer-name;
}
rpf-check;
}
family iso {
    address interface-address;
    mtu bytes;
}
family mpls {
    mtu bytes;
}
}
proxy-arp (restricted | unrestricted);
swap-by-poppush;
(traps | no-traps);
vlan-id (VLAN Tagging and Layer 3 Subinterfaces) vlan-id-number;
}

```

```

    vlan-tagging;
  }
}

```

Unsupported Statements in the [edit interfaces xe] Hierarchy Level

All statements in the **[edit interfaces xe]** hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented.

Related Documentation

- [\[edit interfaces\] Configuration Statement Hierarchy on EX Series Switches on page 141](#)

[edit protocols lacp] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the **[edit protocols lacp]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see *EX Series Switch Software Features Overview*.

This topic lists:

- [Supported Statements in the \[edit protocols lacp\] Hierarchy Level on page 173](#)
- [Unsupported Statements in the \[edit protocols lacp\] Hierarchy Level on page 174](#)

Supported Statements in the [edit protocols lacp] Hierarchy Level

The following hierarchy shows the **[edit protocols lacp]** configuration statements supported on EX Series switches:

```

protocols {
  lacp {
    ppm {
      centralized
    }
    traceoptions {
      file <filename> <files number> <match regular-expression> <size maximum-file-size>
        <world-readable | no-world-readable>;
      flag flag;
      no-remote-trace;
    }
  }
}

```

accounting-profile

Syntax	accounting-profile <i>name</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i>], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit interfaces interface-range <i>name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Enable collection of accounting data for the specified physical or logical interface or interface range.
Options	<i>name</i> —Name of the accounting profile.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Applying an Accounting Profile to the Physical Interface on page 98• Applying an Accounting Profile to the Logical Interface on page 101

address

```

Syntax  address address {
        arp ip-address (mac | multicast-mac) mac-address <publish>;
        broadcast address;
        destination address;
        destination-profile name;
        eui-64;
        master-only;
        multipoint-destination address dlcid dlcid-identifier;
        multipoint-destination address {
            epd-threshold cells;
            inverse-arp;
            oam-liveness {
                up-count cells;
                down-count cells;
            }
            oam-period (disable | seconds);
            shaping {
                (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained rate burst
                 length);
                queue-length number;
            }
            vci vpi-identifier.vci-identifier;
        }
        primary;
        preferred;
        (vrrp-group | vrrp-inet6-group) group-number {
            (accept-data | no-accept-data);
            advertise-interval seconds;
            authentication-type authentication;
            authentication-key key;
            fast-interval milliseconds;
            (preempt | no-preempt) {
                hold-time seconds;
            }
            priority-number number;
            track {
                priority-cost seconds;
                priority-hold-time interface-name {
                    interface priority;
                    bandwidth-threshold bits-per-second {
                        priority;
                    }
                }
            }
            route ip-address/mask routing-instance instance-name priority-cost cost;
        }
        virtual-address [ addresses ];
    }
}

```

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family *family*],
 [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*
 family *family*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 11.1 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Configure the interface address.



NOTE: The vrrp High Availability functionality is not available on the QFX Series.

Options *address*—Address of the interface.

- In Junos OS Release 13.3 and later, when you configure an IPv6 host address and an IPv6 subnet address on an interface, the commit operation fails.
- In releases earlier than Junos OS Release 13.3, when you use the same configuration on an interface, the commit operation succeeds, but only one of the IPv6 addresses that was entered is assigned to the interface. The other address is not applied.



NOTE: If you configure the same address on multiple interfaces in the same routing instance, Junos OS uses only the first configuration. The remaining address configurations are ignored and can leave interfaces without an address. Interfaces that do not have an assigned address cannot be used as a donor interface for an unnumbered Ethernet interface.

For example, in the following configuration the address configuration of interface xe-0/0/1.0 is ignored:

```
interfaces {
  xe-0/0/0 {
    unit 0 {
      family inet {
        address 192.168.1.1/24;
      }
    }
  }
  xe-0/0/1 {
    unit 0 {
      family inet {
        address 192.168.1.1/24;
      }
    }
  }
}
```

For more information on configuring the same address on multiple interfaces, see [“Configuring the Interface Address” on page 84](#).

The remaining statements are explained separately.



NOTE: The `edit logical-systems` hierarchy is not available on QFabric systems.

Required Privilege Level	interface—To view this statement in the configuration.
	interface-control—To add this statement to the configuration.

- Related Documentation**
- *Configuring the Protocol Family*
 - *Junos OS Administration Library for Routing Devices*
 - [family on page 196](#)
 - *negotiate-address*
 - *unnumbered-address (Ethernet)*
 - *Junos OS Administration Library for Routing Devices*

aggregated-devices

Syntax	<pre>aggregated-devices { ethernet (Aggregated Devices) { device-count number; lacp } }</pre>
Hierarchy Level	[edit chassis]
Release Information	Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	<p>Configure properties for aggregated devices on the switch.</p> <p>The remaining statements are explained separately.</p>
Default	Aggregated devices are disabled.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch on page 31 • Configuring Aggregated Ethernet Links (CLI Procedure) on page 109 • Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure) on page 114 • Understanding Aggregated Ethernet Interfaces and LACP on page 8 • <i>Junos OS Ethernet Interfaces Configuration Guide</i>

aggregated-ether-options

Syntax	<pre> aggregated-ether-options { (flow-control no-flow-control); lacp { (active passive); admin-key <i>key</i>; periodic <i>interval</i>; system-id <i>mac-address</i>; } (link-protection no-link-protection); link-protection-sub-group <i>group-name</i> { [primary backup]; } link-speed <i>speed</i>; (loopback no-loopback); minimum-links <i>number</i>; system-priority } </pre>
Hierarchy Level	[edit interfaces <i>aex</i>]
Release Information	Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	<p>Configure the aggregated Ethernet properties of a specific aggregated Ethernet interface.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch on page 31 • Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch on page 37 • Configuring Aggregated Ethernet Links (CLI Procedure) on page 109 • Configuring Aggregated Ethernet LACP (CLI Procedure) on page 113 • Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure) on page 114 • Understanding Aggregated Ethernet Interfaces and LACP on page 8 • Junos OS Ethernet Interfaces Configuration Guide

alarm (optics-options)

Syntax	alarm low-light-alarm { (link-down syslog); }
Hierarchy Level	[edit interfaces <i>interface-name</i> optics-options]
Release Information	Statement introduced in Junos OS Release 10.0. Statement introduced in Junos OS Release 12.1 for EX Series switches.
Description	Specify the action to take if the receiving optics signal is below the optics low-light alarm threshold.
Options	link-down —Drop the 10-Gigabit Ethernet link and marks link as down. syslog —Write the optics information to the system log.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring 10-Gigabit Ethernet Link Down Notification for Optics Options Alarm or Warning</i>• <i>100-Gigabit Ethernet OTN Options Configuration Overview</i>

Default Autonegotiation is automatically enabled. No explicit action is taken after the autonegotiation is complete or if the negotiation fails.

Options **remote-fault (local-interface-online | local-interface-offline)**—(Optional) For M Series, MX Series, T Series, TX Matrix routers, and ACX Series routers only, manually configure remote fault on an interface.

Default: local-interface-online

Required Privilege interface—To view this statement in the configuration.
Level interface-control—To add this statement to the configuration.

Related Documentation

- *Gigabit Ethernet Autonegotiation Overview*
- [Configuring Gigabit Ethernet Interfaces \(CLI Procedure\) on page 64](#)
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*

bandwidth (Interfaces)

Syntax	<code>bandwidth rate;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Configure an informational-only bandwidth value for an interface. This statement is valid for all logical interface types except multilink and aggregated interfaces.




NOTE: We recommend that you be careful when setting this value. Any interface bandwidth value that you configure using the `bandwidth` statement affects how the interface cost is calculated for a dynamic routing protocol, such as OSPF. By default, the interface cost for a dynamic routing protocol is calculated using the following formula:

$$\text{cost} = \text{reference-bandwidth} / \text{bandwidth},$$

where bandwidth is the physical interface speed. However, if you specify a value for bandwidth using the `bandwidth` statement, that value is used to calculate the interface cost, rather than the actual physical interface bandwidth.

Options	rate —Peak rate, in bits per second (bps) or cells per second (cps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation k (1000), m (1,000,000), or g (1,000,000,000). You can also specify a value in cells per second by entering a decimal number followed by the abbreviation c ; values expressed in cells per second are converted to bits per second by means of the formula 1 cps = 384 bps. Range: Not limited.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring the Interface Bandwidth on page 86

broadcast

Syntax	<code>broadcast address;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> address <i>address</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> address <i>address</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Set the broadcast address on the network or subnet. On a subnet you cannot specify a host address of 0 (0.0.0.0), nor can you specify a broadcast address (255.255.255.255). For example, in the statement set interface ge-0/0/0 unit 0 family inet address 10.1.1.0/24 , the subnet address 10.1.1.0 has the host address of 0. Hence, you cannot configure this address. Similarly, for the subnet, you cannot use the broadcast address 10.1.1.255/24.
Default	The default broadcast address has a host portion of all ones.
Options	address —Broadcast address. The address must have a host portion of either all ones or all zeros. You cannot specify the addresses 0.0.0.0 or 255.255.255.255.
<div>  NOTE: The edit logical-systems hierarchy is not available on QFabric systems. </div>	
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> Configuring the Interface Address on page 84

chassis

```
Syntax  chassis {
        aggregated-devices {
            ethernet (Aggregated Devices) {
                device-count number;
            }
        }
        auto-image-upgrade;
        fpc slot {
            pic pic-number {
                sfpplus {
                    pic-mode mode;
                }
            }
            power-budget-priority priority;
        }
        lcd-menu {
            fpc slot-number {
                menu-item (menu-name | menu-option) {
                    disable;
                }
            }
        }
        nssu {
            upgrade-group group-name {
                fpcs (NSSU Upgrade Groups) (slot-number | [list-of-slot-numbers]);
                member (NSSU Upgrade Groups) member-id {
                    fpcs (NSSU Upgrade Groups) (slot-number | [list-of-slot-numbers]);
                }
            }
        }
        psu {
            redundancy {
                n-plus-n (Power Management);
            }
        }
        redundancy {
            graceful-switchover;
        }
    }
```

Hierarchy Level [edit]

Release Information Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Configure chassis-specific properties for the switch.

The remaining statements are explained separately.

Required Privilege interface—To view this statement in the configuration.

Level interface-control—To add this statement to the configuration.

- Related Documentation**
- [Configuring Aggregated Ethernet Links \(CLI Procedure\) on page 109](#)
 - [Upgrading Software by Using Automatic Software Download](#)
 - [Configuring the LCD Panel on EX Series Switches \(CLI Procedure\)](#)
 - [Configuring Graceful Routing Engine Switchover in a Virtual Chassis \(CLI Procedure\)](#)
 - [Configuring Power Supply Redundancy \(CLI Procedure\)](#)
 - [Configuring the Power Priority of Line Cards \(CLI Procedure\)](#)
 - [Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade \(CLI Procedure\)](#)

description

Syntax	description text;
Hierarchy Level	[edit interfaces <i>interface-name</i>], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	<p>Provide a textual description of the interface or the logical unit. Any descriptive text you include is displayed in the output of the show interfaces commands, and is also exposed in the ifAlias Management Information Base (MIB) object. It has no effect on the operation of the interface on the router or switch.</p> <p>The textual description can also be included in the extended DHCP relay option 82 Agent Circuit ID suboption.</p>
Options	text —Text to describe the interface. If the text includes spaces, enclose the entire text in quotation marks.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Adding an Interface Description to the Configuration• Adding a Logical Unit Description to the Configuration on page 79• Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 64• Using DHCP Relay Agent Option 82 Information

destination (Tunnels)

Syntax	<code>destination address;</code>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i>],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet unnumbered-address <i>interface-name</i>],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> tunnel],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet unnumbered-address <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> tunnel]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.1 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	For encrypted, PPP-encapsulated, and tunnel interfaces, specify the remote address of the connection.
Options	address —Address of the remote side of the connection.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring the Interface Address on page 84 • Configuring Generic Routing Encapsulation Tunneling (CLI Procedure) on page 130 • <i>Junos OS Services Interfaces Library for Routing Devices</i> • <i>point-to-point</i>

device-count

Syntax	<code>device-count <i>number</i>;</code>
Hierarchy Level	[edit chassis aggregated-devices ethernet (Aggregated Devices)]
Release Information	Statement introduced in Junos OS Release 9.0 for EX Series switches. Range updated in Junos OS Release 9.5 for EX Series switches.
Description	Configure the number of aggregated Ethernet logical devices available to the switch.
Options	<p><i>number</i>—Maximum number of aggregated Ethernet logical interfaces on the switch.</p> <p>Range: 1 through 32 for EX2200, EX3200, and standalone EX3300 switches and for EX3300 Virtual Chassis</p> <p>Range: 1 through 64 for standalone EX4200, standalone EX4500, and EX6200 switches and for EX4200 and EX4500 Virtual Chassis</p> <p>Range: 1 through 239 for EX8200 Virtual Chassis</p> <p>Range: 1 through 255 for standalone EX8200 switches</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch on page 31• Configuring Aggregated Ethernet Links (CLI Procedure) on page 109• Junos OS Network Interfaces Configuration Guide

eui-64

Syntax	eui-64;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>number</i> family inet6 address <i>address</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.3 for EX Series switches. Statement introduced in Junos OS Release 12.2 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	For interfaces that carry IP version 6 (IPv6) traffic, automatically generate the host number portion of interface addresses.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Interface Address on page 84

- [Configuring LACP Link Protection of Aggregated Ethernet Interfaces \(CLI Procedure\) on page 114](#)
- [Understanding Aggregated Ethernet Interfaces and LACP on page 8](#)
- [EX Series Switches Interfaces Overview on page 3](#)
- [Junos OS Ethernet Interfaces Configuration Guide](#)

ethernet (Aggregated Devices)

Syntax	<pre>ethernet { device-count <i>number</i>; lacp { link-protection { non-revertive; } system-priority; } }</pre>
Hierarchy Level	[edit chassis aggregated-devices]
Release Information	Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	<p>Configure properties for Ethernet aggregated devices on the switch.</p> <p>The remaining statement is explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Aggregated Ethernet Links (CLI Procedure) on page 109 • Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure) on page 114 • Junos OS Ethernet Interfaces Configuration Guide


family (for EX Series switches)

Syntax	family ccc on page 196 family ethernet-switching on page 196 family inet on page 196 family inet6 on page 197 family iso on page 197 family mpls on page 197
family ccc	family ccc;
family ethernet-switching	<pre> family ethernet-switching { filter [input output] <i>filter-name</i>; native-vlan-id <i>vlan-id</i>; port-mode <i>mode</i>; vlan (802.1Q Tagging) { members [(all <i>names</i> <i>vlan-ids</i>)]; } }</pre>
family inet	<pre> family inet { address <i>address</i> { arp <i>ip-address</i> (mac multicast-mac) <i>mac-address</i> <publish>; broadcast; preferred; primary; vrrp-group <i>group-id</i> { advertise-interval <i>milliseconds</i>; preempt no-preempt { hold-time <i>seconds</i>; } priority <i>number</i>; virtual-address [<i>addresses</i>]; virtual-link-local-address <i>ip-address</i>; } } dhcp { client-identifier (ascii <i>ascii</i> hexadecimal <i>hexadecimal</i>); lease-time (<i>seconds</i> infinite); retransmission-attempt <i>number</i>; retransmission-interval <i>seconds</i>; server-address <i>ip-address</i>; update-server; vendor-id <i>vendor-id</i>; } filter { input <i>filter-name</i>; output <i>filter-name</i>; } mtu <i>bytes</i>; no-redirects; no-neighbor-learn; primary; rpf-check; }</pre>

	<pre> targeted-broadcast; } </pre>
family inet6	<pre> family inet6 { address address { eui-64; nd6-stale-time seconds; ndp ip-address (mac multicast-mac) mac-address <publish>; preferred; primary; vrrp-inet6-group group-id { inet6-advertise-interval milliseconds; preempt preempt { hold-time seconds; } priority number; virtual-inet6-address [addresses]; virtual-link-local-address ipv6-address; } } (dad-disable no-dad-disable); filter { input filter-name; output filter-name; } mtu bytes; no-neighbor-learn; rpf-check; } </pre>
family iso	<pre> family iso { address interface-address; mtu bytes; } </pre>
family mpls	<pre> family mpls { mtu bytes; } </pre>
Hierarchy Level	<pre> [edit interfaces interface-name unit logical-unit-number], [edit interfaces interface-range name unit logical-unit-number] </pre>
Release Information	<p>Statement introduced in Junos OS Release 9.0 for EX Series switches, including options ethernet-switching, inet, and iso.</p> <p>Option inet6 introduced in Junos OS Release 9.3 for EX Series switches.</p> <p>Options ccc and mpls introduced in Junos OS Release 9.5 for EX Series switches.</p>
Description	<p>Configure protocol family information for the logical interface on the switch.</p> <p>You must configure a logical interface to be able to use the physical device.</p>


Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring a DHCP Server on Switches (CLI Procedure)</i>• <i>Example: Configuring MPLS on EX8200 and EX4500 Switches</i>• Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 64• Configuring Aggregated Ethernet Links (CLI Procedure) on page 109• <i>Configuring Routed VLAN Interfaces (CLI Procedure)</i>

filter

Syntax	<pre>filter { group <i>filter-group-number</i>; input <i>filter-name</i>; input-list [<i>filter-names</i>]; output <i>filter-name</i>; output-list [<i>filter-names</i>]; }</pre>
Hierarchy Level	<pre>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>]</pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>
Description	<p> NOTE: On EX Series switches, the <code>group</code>, <code>input-list</code>, <code>output-filter</code> statements are not supported under the <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet]</code>, <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6]</code>, and <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family mpls]</code> hierarchies.</p> <p>Apply a filter to an interface. You can also use filters for encrypted traffic. When you configure filters, you can configure them under the family ethernet-switching, inet, inet6, mpls, or vpls only.</p>
Options	<p>group <i>filter-group-number</i>—Define an interface to be part of a filter group. The default filter group number is 0.</p> <p>Range: 0 through 255</p> <p>input <i>filter-name</i>—Name of one filter to evaluate when packets are received on the interface.</p> <p>output <i>filter-name</i>—Name of one filter to evaluate when packets are transmitted on the interface.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Applying a Filter to an Interface</i> • <i>Junos OS Administration Library for Routing Devices</i> • Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 64 • <i>Configuring Firewall Filters (CLI Procedure)</i>

- [family on page 196](#)

flow-control

Syntax	(flow-control no-flow-control);
Hierarchy Level	[edit interfaces <i>interface-name</i> aggregated-ether-options], [edit interfaces <i>interface-name</i> ether-options], [edit interfaces <i>interface-name</i> fastether-options], [edit interfaces <i>interface-name</i> ggether-options], [edit interfaces <i>interface-name</i> multiservice-options], [edit interfaces interface-range <i>name</i> aggregated-ether-options], [edit interfaces interface-range <i>name</i> ether-options]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 in EX Series switches. Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.
Description	For aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet interfaces only, explicitly enable flow control, which regulates the flow of packets from the router or switch to the remote side of the connection. Enabling flow control is useful when the remote device is a Gigabit Ethernet switch. Flow control is not supported on the 4-port Fast Ethernet PIC.
<div>  <p>NOTE: On the Type 5 FPC, to prioritize control packets in case of ingress oversubscription, you must ensure that the neighboring peers support MAC flow control. If the peers do not support MAC flow control, then you must disable flow control.</p> </div>	
Default	Flow control is enabled.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Flow Control on page 82 • Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 64 • Configuring Gigabit Ethernet Interfaces (CLI Procedure)

force-up

Syntax	force-up;
Hierarchy Level	[edit interfaces <i>interface-name</i> ether-options 802.3ad lacp]
Release Information	Statement introduced in Junos OS Release 10.0 for EX Series switches.
Description	Set the state of the interface as UP when the peer has limited LACP capability.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 64• Configuring Gigabit Ethernet Interfaces (CLI Procedure)• Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 68• Understanding Aggregated Ethernet Interfaces and LACP on page 8• Junos OS Ethernet Interfaces Configuration Guide

gratuitous-arp-reply

Syntax	(gratuitous-arp-reply no-gratuitous-arp-reply);
Hierarchy Level	[edit interfaces <i>interface-name</i>] [edit interfaces <i>interface-range</i> <i>interface-range-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	For Ethernet interfaces, enable updating of the Address Resolution Protocol (ARP) cache for gratuitous ARPs.
Default	Updating of the ARP cache is disabled on all Ethernet interfaces.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Interfaces Overview

Range: 0 through 4,294,967,295

Default: 0 (interface transitions are not damped)

up *milliseconds*—Hold time to use when an interface transitions from down to up. Junos OS advertises the transition within 100 milliseconds of the time value you specify.

Range: 0 through 4,294,967,295

Default: 0 (interface transitions are not damped)

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *advertise-interval*
- [interfaces \(for EX Series switches\) on page 207](#)
- *Physical Interface Damping Overview*
- [Damping Shorter Physical Interface Transitions on page 135](#)
- *Damping Longer Physical Interface Transitions*

ieee-802-3az-eee

Syntax ieee-802-3az-eee;

Hierarchy Level [edit interfaces *interface-name* ether-options]

Release Information Statement introduced in Junos OS Release 12.2 for EX Series switches.

Description Configure Energy Efficient Ethernet (EEE) on an EEE-capable Base-T copper interface.

Default EEE is disabled on EEE-capable interfaces.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- [Configuring Energy Efficient Ethernet on Interfaces \(CLI Procedure\) on page 133](#)

interface-range

```
Syntax interface-range name {
    accounting-profile name;
    description text;
    disable;
    ether-options {
        802.3ad {
            aex;
            (backup | primary);
            lacp {
                force-up;
            }
        }
        (auto-negotiation | no-auto-negotiation);
        (flow-control | no-flow-control);
        ieee-802-3az-eee;
        link-mode mode;
        (loopback | no-loopback);
        speed (auto-negotiation | speed);
    }
    (gratuitous-arp-reply | no-gratuitous-arp-reply);
    hold-time up milliseconds down milliseconds;
    member interface-name;
    member-range starting-interface name to ending-interface name;
    mtu bytes;
    no-gratuitous-arp-request;
    traceoptions {
        flag flag;
    }
    (traps | no-traps);
    unit logical-unit-number {
        accounting-profile name;
        bandwidth rate;
        description text;
        disable;
        family family-name {...}
        proxy-arp (restricted | unrestricted);
        (traps | no-traps);
        vlan-id (VLAN Tagging and Layer 3 Subinterfaces) vlan-id-number;
    }
    vlan-tagging;
}
```

Hierarchy Level [edit [interfaces](#)]

Release Information Statement introduced in Junos OS Release 10.0 for EX Series switches.

Description Group interfaces that share a common configuration profile.



NOTE: You can specify interface ranges only for Gigabit and 10-Gigabit Ethernet interfaces.

Options *name*—Name of the interface range.



NOTE: You can use regular expressions and wildcards to specify the interfaces in the member configuration. Do not use wildcards for interface types.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

Related Documentation

- [Configuring Gigabit Ethernet Interfaces \(CLI Procedure\) on page 64](#)
- [Configuring Gigabit Ethernet Interfaces \(CLI Procedure\)](#)
- [Understanding Interface Ranges on EX Series Switches on page 17](#)
- [Understanding Interface Ranges on EX Series Switches](#)
- [EX Series Switches Interfaces Overview on page 3](#)
- [Junos OS Interfaces Fundamentals Configuration Guide](#)

interfaces (for EX Series switches)

Syntax [interfaces ae on page 207](#)
[interfaces ge on page 207](#)
[interfaces interface-range on page 209](#)
[interfaces lo0 on page 209](#)
[interfaces me0 on page 210](#)
[interfaces traceoptions on page 210](#)
[interfaces vlan on page 210](#)
[interfaces vme on page 211](#)
[interfaces xe on page 212](#)

```

interfaces ae  aex {
    accounting-profile name;
    aggregated-ether-options {
        (flow-control | no-flow-control);
        lacp {
            (active | passive);
            admin-key key;
            periodic interval;
            system-id mac-address;
        }
        (link-protection | no-link-protection);
        link-speed speed;
        (loopback | no-loopback);
        minimum-links number;
    }
    description text;
    disable;
    (gratuitous-arp-reply | no-gratuitous-arp-reply);
    mtu bytes;
    no-gratuitous-arp-request;
    traceoptions {
        flag flag;
    }
    (traps | no-traps);
    unit logical-unit-number {
        accounting-profile name;
        bandwidth rate;
        description text;
        disable;
        family family-name {...}
        proxy-arp (restricted | unrestricted);
        (traps | no-traps);
        vlan-id vlan-id-number;
    }
    vlan-tagging;
}

interfaces ge  ge-fpc/pic/port {
    accounting-profile name;
    description text;
    disable;
    ether-options {
        802.3ad {

```

```
    aex;
    (backup | primary);
    lacp {
        force-up;
    }
}
(auto-negotiation | no-auto-negotiation);
(flow-control | no-flow-control);
ieee-802-3az-eee;
link-mode mode;
(loopback | no-loopback);
speed (auto-negotiation | speed);
}
(gratuitous-arp-reply | no-gratuitous-arp-reply);
hold-time up milliseconds down milliseconds;
media-type;
mtu bytes;
no-gratuitous-arp-request;
traceoptions {
    flag flag;
}
(traps | no-traps);
unit logical-unit-number {
    accounting-profile name;
    bandwidth rate;
    description text;
    disable;
    family family-name {...}
    proxy-arp (restricted | unrestricted);
    (traps | no-traps);
    vlan-id vlan-id-number;
}
vlan-tagging;
}
```

```

interfaces interface-range name {
    accounting-profile name;
    description text;
    disable;
    ether-options {
        802.3ad {
            aex;
            (backup | primary);
            lACP {
                force-up;
            }
        }
        (auto-negotiation | no-auto-negotiation);
        (flow-control | no-flow-control);
        ieee-802-3az-eee;
        link-mode mode;
        (loopback | no-loopback);
        speed (auto-negotiation | speed);
    }
    (gratuitous-arp-reply | no-gratuitous-arp-reply);
    hold-time up milliseconds down milliseconds;
    member interface-name;
    member-range starting-interface name to ending-interface name;
    mtu bytes;
    unit logical-unit-number {
        accounting-profile name;
        bandwidth rate;
        description text;
        disable;
        family family-name {...}
        proxy-arp (restricted | unrestricted);
        (traps | no-traps);
        vlan-id vlan-id-number;
    }
    vlan-tagging;
}

interfaces lo0 lo0 {
    accounting-profile name;
    description text;
    disable;
    hold-time up milliseconds down milliseconds;
    traceoptions {
        flag flag;
    }
    (traps | no-traps);
    unit logical-unit-number {
        accounting-profile name;
        bandwidth rate;
        description text;
        disable;
        family family-name {...}
        (traps | no-traps);
    }
}

```

```

interfaces me0 me0 {
    accounting-profile name;
    description text;
    disable;
    (gratuitous-arp-reply | no-gratuitous-arp-reply);
    hold-time up milliseconds down milliseconds;
    no-gratuitous-arp-request;
    traceoptions {
        flag flag;
    }
    (traps | no-traps);
    unit logical-unit-number {
        accounting-profile name;
        bandwidth rate;
        description text;
        disable;
        family family-name {...}
        (traps | no-traps);
        vlan-id vlan-id-number;
    }
    vlan-tagging;
}

interfaces traceoptions traceoptions {
    file <filename> <files number> <match regular-expression> <size size> <world-readable |
    no-world-readable>;
    flag flag <disable>;
    no-remote-trace;
}

interfaces vlan vlan {
    accounting-profile name;
    description text;
    disable;
    (gratuitous-arp-reply | no-gratuitous-arp-reply);
    hold-time up milliseconds down milliseconds;
    mtu bytes;
    no-gratuitous-arp-request;
    traceoptions {
        flag flag;
    }
    (traps | no-traps);
    unit logical-unit-number {
        accounting-profile name;
        bandwidth rate;
        description text;
        disable;
        family family-name {...}
        proxy-arp (restricted | unrestricted);
        (traps | no-traps);
    }
}

```

```
interfaces vme    vme {  
    accounting-profile name;  
    description text;  
    disable;  
    (gratuitous-arp-reply | no-gratuitous-arp-reply);  
    hold-time up milliseconds down milliseconds;  
    mtu bytes;  
    no-gratuitous-arp-request;  
    traceoptions {  
        flag flag;  
    }  
    (traps | no-traps);  
    unit logical-unit-number {  
        accounting-profile name;  
        bandwidth rate;  
        description text;  
        disable;  
        family family-name {...}  
        (traps | no-traps);  
        vlan-id vlan-id-number;  
    }  
    vlan-tagging;  
}
```


**Related
Documentation**

- [Configuring Gigabit Ethernet Interfaces \(CLI Procedure\) on page 64](#)
- [Configuring Aggregated Ethernet Links \(CLI Procedure\) on page 109](#)
- [Configuring a Layer 3 Subinterface \(CLI Procedure\) on page 123](#)
- [Configuring Routed VLAN Interfaces \(CLI Procedure\)](#)
- [Configuring the Virtual Management Ethernet Interface for Global Management of an EX Series Virtual Chassis \(CLI Procedure\)](#)
- [EX Series Switches Interfaces Overview on page 3](#)
- [Junos OS Interfaces Fundamentals Configuration Guide](#)
- [Junos OS Ethernet Interfaces Configuration Guide](#)

admin-key *number*—Specify an administrative key for the router or switch.



NOTE: You must also configure multichassis link aggregation (MC-LAG) when you configure the **admin-key**.

fast-failover—Specify to override the IEEE 802.3ad standard and allow the standby link to receive traffic. Overriding the default behavior facilitates subsecond failover.

passive—Respond to LACP packets.

The remaining statements are explained separately.

Required Privilege Level	interface—To view this statement in the configuration.
	interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring LACP for Aggregated Ethernet Interfaces</i>

link-protection-sub-group (802.3ad)

Syntax	link-protection-sub-group <i>group-name</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> ether-options 802.3ad]
Release Information	Statement introduced in Junos OS Release 14.1X53-D10 for EX Series switches.
Description	<p>Add an interface in an aggregated ethernet bundle into a link protection subgroup.</p> <p>A link protection subgroup is created and named using the link-protection-sub-group statement in the [edit interfaces aex aggregated-ether-options] hierarchy.</p>
Options	<p><i>group-name</i>—Name of the link protection subgroup that will include this interface after this statement is entered. The link protection subgroup is named when it is created using the link-protection-sub-group statement in the [edit interfaces aex aggregated-ether-options] hierarchy.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring Aggregated Ethernet Link Protection on page 118• Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure) on page 114

link-protection-sub-group (aggregated-ether-options)

Syntax	link-protection-sub-group <i>group-name</i> { [primary backup]; }
Hierarchy Level	[edit interfaces aex aggregated-ether-options]
Release Information	Statement introduced in Junos OS Release 14.1X53-D10 for EX Series switches.
Description	<p>Create and name a link protection subgroup.</p> <p>Link protection subgroups allow you to provide link protection to a collection of Ethernet links within a LAG bundle. If you need to provide link protection to a single link in an aggregated ethernet bundle, you do not need to configure link protection subgroups.</p> <p>A link protection subgroup includes multiple links within the aggregated ethernet bundle. If one link in the primary link protection subgroup fails, traffic is forwarded over the links in the backup link protection subgroup.</p> <p>Links within the aggregated ethernet bundle are added to the link protection subgroup using the link-protection-sub-group statement in the [edit interfaces <i>interface-name</i> ether-options 802.3ad] hierarchy.</p>
Options	<p><i>group-name</i>—Creates and names the link protection subgroup. The name is created by the user.</p> <p>primary—Specifies that the subgroup is the primary subgroup.</p> <p>backup—Specifies that the subgroup is the backup subgroup.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Aggregated Ethernet Link Protection on page 118 • Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure) on page 114

- **mixed**—Links are of various speeds.
- **oc192**—Links are OC192.

mixed—Enables bundling of different Ethernet rate links in the same Aggregated Ethernet interface on PTX Series routers.

Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
---------------------------------	---

Related Documentation	<ul style="list-style-type: none">• <i>Aggregated Ethernet Interfaces Overview</i>• Configuring Aggregated Ethernet Link Speed on page 120• <i>Configuring Mixed Rates and Mixed Modes on Aggregated Ethernet Bundles</i>• Configuring Aggregated Ethernet Links (CLI Procedure) on page 109• Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch on page 31
------------------------------	--

member (Interface Ranges)

Syntax	<code>member <i>interface-name</i>;</code>
Hierarchy Level	[edit interfaces interface-range <i>interface-range-name</i>]
Release Information	Statement introduced in Junos OS Release 10.0 for EX Series switches.
Description	Specify the name of the member interface belonging to an interface range on the EX Series switch.
Options	<i>interface-name</i> —Name of the interface.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 64 • Configuring Gigabit Ethernet Interfaces (CLI Procedure) • Understanding Interface Ranges on EX Series Switches on page 17 • Understanding Interface Ranges on EX Series Switches • EX Series Switches Interfaces Overview on page 3 • Junos OS Interfaces Fundamentals Configuration Guide

member-range

Syntax	<code>member-range <i>starting-interface-name</i> to <i>ending-interface-name</i>;</code>
Hierarchy Level	[edit interfaces interface-range interface-range-name]
Release Information	Statement introduced in Junos OS Release 10.0 for EX Series switches.
Description	Specify the names of the first and last members of a sequence of interfaces belonging to an interface range.
Options	Range: <i>Starting interface-name</i> to <i>ending interface-name</i> —The name of the first member and the name of the last member in the interface sequence.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 64• Configuring Gigabit Ethernet Interfaces (CLI Procedure)• Understanding Interface Ranges on EX Series Switches on page 17• Understanding Interface Ranges on EX Series Switches• EX Series Switches Interfaces Overview on page 3• Junos OS Interfaces Fundamentals Configuration Guide



NOTE:

The MTU for an IRB interface is calculated by removing the Ethernet header overhead [6(DMAC)+6(SMAC)+2(EtherType)]. Because, the MTU is the lower value of the MTU configured on the IRB interface and the MTU configured on the IRB's associated bridge domain IFDs or IFLs, the IRB MTU is calculated as follows:

- In case of Layer 2 IFL configured with the `flexible-vlan-tagging` statement, the IRB MTU is calculated by including 8 bytes overhead (SVLAN+CVLAN).
 - In case of Layer 2 IFL configured with the `vlan-tagging` statement, the IRB MTU is calculated by including a single VLAN 4 bytes overhead.
-

non-revertive (Interfaces)

Syntax	non-revertive;
Hierarchy Level	[edit interfaces aeX aggregated-ether-options lacp link-protection]
Release Information	Statement introduced in Junos OS Release 9.3. Statement introduced in Junos OS Release 11.4 for EX Series switches.
Description	Disable the ability to switch to a better priority link (if one is available) once a link is established as active and collection distribution is enabled.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• link-protection on page 219• Configuring Aggregated Ethernet Link Protection on page 118• Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure) on page 114

reflective-relay

Syntax	reflective-relay;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family ethernet-switching]
Release Information	Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D35 for the EX Series.
Description	Configure a switch interface to return packets back to a device on the same interface that was used to deliver the packets.
Default	Switch interfaces are not configured for reflective relay.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring Reflective Relay for Use with VEPA Technology</i>• Configuring Reflective Relay on page 136

source

Syntax	<code>source <i>source-address</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> tunnel]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.1 for EX Series switches. Statement introduced in Junos OS Release 13.2 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Specify the source address of the tunnel.
Default	If you do not specify a source address, the tunnel uses the unit's primary address as the source address of the tunnel.
Options	<i>source-address</i> —Address of the local side of the tunnel. This is the address that is placed in the outer IP header's source field.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Tunnel Services Overview</i>

system-priority

Syntax	<code>system-priority <i>priority</i>;</code>
Hierarchy Level	[edit interfaces aeX aggregated-ether-options lacp]
Release Information	Statement introduced in Junos OS Release 9.3. Statement introduced in Junos OS Release 11.4 for EX Series switches.
Description	<p>Define LACP system priority at the aggregated Ethernet interface level. This system priority value takes precedence over a system priority value configured at the global [edit chassis] hierarchy level.</p> <p>The device with the lower system priority value determines which links between LACP partner devices are active and which are in standby for each LACP group. The device on the controlling end of the link uses port priorities to determine which ports are bundled into the aggregated bundle and which ports are put in standby mode. Port priorities on the other device (the noncontrolling end of the link) are ignored. In priority comparisons, numerically lower values have higher priority. Therefore, the system with the numerically lower value (higher priority value) for LACP system priority becomes the controlling system. If both devices have the same LACP system priority (for example, they are both configured with the default setting of 127), the device MAC address determines which switch is in control.</p>
Options	<p><i>priority</i>—Priority for the aggregated Ethernet system. A smaller value indicates a higher priority.</p> <p>Range: 0 through 65535</p> <p>Default: 127</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

targeted-broadcast

Syntax	targeted-broadcast;
Hierarchy Level	[edit interfaces ge-chassis/slot/port unit logical-unit-number family inet]
Release Information	Statement introduced in Junos OS Release 9.4 for EX Series switches.
Description	Enable IP directed broadcast on a specified subnet.
Default	IP directed broadcast is disabled.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring IP Directed Broadcast on a Switch on page 54• Configuring IP Directed Broadcast (CLI Procedure) on page 126• Understanding IP Directed Broadcast on page 16

- **media**—Interface media changes
- **q921**—Trace ISDN Q.921 frames
- **q931**—Trace ISDN Q.931 frames

Required Privilege interface—To view this statement in the configuration.
Level interface-control—To add this statement to the configuration.

Related Documentation

- *Tracing Operations of an Individual Router Interface*

tunnel-port

Syntax	<code>tunnel-port <i>port-number</i> tunnel-services;</code>
Hierarchy Level	[edit chassis fpc slot <i>pic</i> <i>pic-number</i>]
Release Information	Statement introduced in Junos OS Release 12.1 for EX Series switches. Statement introduced in Junos OS Release 13.2 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure the port number for generic routing encapsulation (GRE) tunneling.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Generic Routing Encapsulation Tunneling (CLI Procedure) on page 130

vlan (802.1Q Tagging)

Syntax	<code>vlan { members [(all names vlan-ids)]; }</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family ethernet-switching]
Release Information	Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	<p>Bind an 802.1Q VLAN tag ID to a logical interface.</p> <p>The remaining statement is explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• <i>show ethernet-switching interfaces</i>• <i>show ethernet-switching interface</i>• <i>Example: Setting Up Bridging with Multiple VLANs for EX Series Switches</i>• <i>Configuring Routed VLAN Interfaces (CLI Procedure)</i>• <i>Configuring Integrated Routing and Bridging Interfaces (CLI Procedure)</i>• <i>Understanding Bridging and VLANs on EX Series Switches</i>• Junos OS Ethernet Interfaces Configuration Guide

vlan-tagging

Syntax	vlan-tagging;
Hierarchy Level	[edit interfaces <i>interface-name</i>]
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.</p> <p>Statement introduced in Junos OS Release 13.2 for PTX Series Routers.</p> <p>Statement introduced in Junos OS Release 14.1X53-D10 for the QFX Series.</p>
Description	For Fast Ethernet and Gigabit Ethernet interfaces and aggregated Ethernet interfaces configured for VPLS, enable the reception and transmission of 802.1Q VLAN-tagged frames on the interface.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• 802.1Q VLANs Overview on page 19• vlan-id on page 267• Configuring a Layer 3 Subinterface (CLI Procedure) on page 123• Configuring Tagged Aggregated Ethernet Interfaces on page 122• Example: Configuring Layer 3 Subinterfaces for a Distribution Switch and an Access Switch on page 42

warning

Syntax	<pre>warning low-light-warning { (link-down syslog); }</pre>
Hierarchy Level	[edit interfaces <i>interface-name</i> optics-options]
Release Information	<p>Statement introduced in Junos OS Release 10.0.</p> <p>Statement introduced in Junos OS Release 12.1 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2 for PTX Series routers.</p>
Description	Specifies the action to take if the receiving optics signal is below the optics low-light warning threshold.
Options	<p>link-down—Drop the 10-Gigabit Ethernet link and marks link as down.</p> <p>syslog—Write the optics information to the system log.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring 10-Gigabit Ethernet Link Down Notification for Optics Options Alarm or Warning</i> • optics-options on page 241 • <i>100-Gigabit Ethernet OTN Options Configuration Overview</i>

PART 3

Administration

- [Routine Monitoring on page 277](#)
- [Operational Commands on page 287](#)

CHAPTER 5

Routine Monitoring

- [Monitoring Interface Status and Traffic on page 277](#)
- [Verifying the Status of a LAG Interface on page 279](#)
- [Verifying That LACP Is Configured Correctly and Bundle Members Are Exchanging LACP Protocol Packets on page 279](#)
- [Verifying That Layer 3 Subinterfaces Are Working on page 281](#)
- [Verifying Unicast RPF Status on page 281](#)
- [Verifying IP Directed Broadcast Status on page 284](#)
- [Verifying That Generic Routing Encapsulation Tunneling Is Working Correctly on page 284](#)
- [Verifying That EEE Is Saving Energy on Configured Ports on page 284](#)

Monitoring Interface Status and Traffic

Purpose



NOTE: This topic applies only to the J-Web Application package.

Use the monitoring functionality to view interface status or to monitor interface bandwidth utilization and traffic statistics on the EX Series switches.

The J-Web interface monitors interface bandwidth utilization and plots real-time charts to display input and output rates in bytes per second. In addition, the Interface monitoring page displays input and output packet counters and error counters in the form of charts.

Alternatively, you can enter the **show** commands in the CLI to view interface status and traffic statistics.



NOTE: For logical interfaces on EX Series switches, the traffic statistics fields in **show interfaces** commands show only control traffic; the traffic statistics do not include data traffic.



NOTE: EX Series switches do not support the collection and reporting of IPv6 transit statistics. Therefore, the IPv6 transit statistics field in the `show interfaces` commands displays all values as 0.

Action To view general interface information in the J-Web interface such as available interfaces, select **Monitor > Interfaces**. Click any interface to view details about its status.

To set up interface monitoring for Virtual Chassis and EX8200 switches, select a member from the **Port for Member** list. Details such as the admin status and link status are displayed in the table. For an EX8200 Virtual Chassis setup, select the member, **FPC**, and the required interface.



NOTE: By default, the details of the first member in the FPC list is displayed. In an EX8200 Virtual Chassis setup, details of the first member and the first FPC is displayed.

You have the following options:

- **Start/Stop**—Starts or stops monitoring the selected interface.
- **Show Graph**—Displays input and output packet counters and error counters in the form of charts. Click the pop-up icon to view the graph in a separate window.
- **Details**—Displays interface information such as general details, traffic statistics, I/O errors, CoS counters, and Ethernet statistics.
- **Refresh Interval (sec)**—Displays the time interval you have set for page refresh.
- **Clear Statistics**—Clears the statistics for the interface selected from the table.

Using the CLI:

- To view interface status for all the interfaces, enter `show interfaces xe-`.
- To view status and statistics for a specific interface, enter `show interfaces xe-interface-name`.
- To view status and traffic statistics for all interfaces, enter either `show interfaces xe-detail` or `show interfaces xe- extensive`.

Meaning In the J-Web interface the charts displayed are:

- **Bar charts**—Display the input and output error counters.
- **Pie charts**—Display the number of broadcast, unicast, and multicast packet counters.

For details about output from the CLI commands, see `show interfaces ge-` (Gigabit Ethernet) or `show interfaces xe-` (10-Gigabit Ethernet).

xe-0/1/0 Defaulted Fast periodic Detached

Meaning This output shows that LACP has been configured with one side as active and the other as passive. When LACP is enabled, at least one side must be set as active for the bundled link to be up.

Verifying That LACP Packets Are Being Exchanged

Purpose Verify that LACP packets are being exchanged between interfaces.

Action Use the **show interfaces aex statistics** command to display LACP BPDU exchange information.

show interfaces ae0 statistics

```
Physical interface: ae0, Enabled, Physical link is Down
  Interface index: 153, SNMP ifIndex: 30
  Link-level type: Ethernet, MTU: 1514, Speed: Unspecified, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Disabled, Minimum links needed: 1,
  Minimum bandwidth needed: 0
  Device flags   : Present Running
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
  Current address: 02:19:e2:50:45:e0, Hardware address: 02:19:e2:50:45:e0
  Last flapped   : Never
  Statistics last cleared: Never
    Input packets : 0
    Output packets: 0
  Input errors: 0, Output errors: 0

Logical interface ae0.0 (Index 71) (SNMP ifIndex 34)
  Flags: Hardware-Down Device-Down SNMP-Traps Encapsulation: ENET2
  Statistics
    Packets      pps      Bytes      bps
  Bundle:
    Input :      0        0          0        0
    Output:      0        0          0        0
  Protocol inet,
    Flags: None
    Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
      Destination: 10.10.10/24, Local: 10.10.10.1, Broadcast: 10.10.10.255
```

Meaning The output here shows that the link is down and that no PDUs are being exchanged (when there is no other traffic flowing on the link).

Related Documentation

- [Configuring Aggregated Ethernet LACP](#)
- [Configuring Aggregated Ethernet LACP \(CLI Procedure\) on page 113](#)
- [Verifying the Status of a LAG Interface](#)
- [Verifying the Status of a LAG Interface on page 279](#)


```

CoS queues      : 8 supported, 8 maximum usable queues
Hold-times     : Up 0 ms, Down 0 ms
Current address: 00:19:e2:50:95:ab, Hardware address: 00:19:e2:50:95:ab
Last flapped   : Never
Statistics last cleared: Never
Traffic statistics:
  Input bytes : 0 0 bps
  Output bytes : 0 0 bps
  Input packets: 0 0 pps
  Output packets: 0 0 pps
IPv6 transit statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runt: 0, Policed discards: 0,
  L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
  FIFO errors: 0, Resource errors: 0
Output errors:
  Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
  FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 8 supported, 4 in use
Queue counters:
  Queued packets  Transmitted packets  Dropped packets

  0 best-effort 0 0 0
  1 assured-fow 0 0 0
  5 expedited-fo 0 0 0
  7 network-cont 0 0 0

Active alarms : LINK
Active defects : LINK
MAC statistics:
  Receive Transmit
  Total octets 0 0
  Total packets 0 0
  Unicast packets 0 0
  Broadcast packets 0 0
  Multicast packets 0 0
  CRC/Align errors 0 0
  FIFO errors 0 0
  MAC control frames 0 0
  MAC pause frames 0 0
  Oversized frames 0
  Jabber frames 0
  Fragment frames 0
  VLAN tagged frames 0
  Code violations 0
Filter statistics:
  Input packet count 0
  Input packet rejects 0
  Input DA rejects 0
  Input SA rejects 0
  Output packet count 0
  Output packet pad count 0
  Output packet error count 0
  CAM destination filters: 0, CAM source filters: 0
Autonegotiation information:

```

```

Negotiation status: Incomplete
Packet Forwarding Engine configuration:
  Destination slot: 1

Logical interface ge-1/0/10.0 (Index 69) (SNMP ifIndex 59) (Generation 135)
Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
Traffic statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
IPv6 transit statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Local statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Transit statistics:
  Input bytes : 0 0 bps
  Output bytes : 0 0 bps
  Input packets: 0 0 pps
  Output packets: 0 0 pps
IPv6 transit statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
  Protocol inet, Generation: 144, Route table: 0
Flags: uRPF
Addresses, Flags: Is-Preferred Is-Primary

```

Meaning The `show interfaces ge-1/0/10 extensive` command (and the `show interfaces ge-1/0/10 detail` command) displays in-depth information about the interface. The **Flags:** output field near the bottom of the display reports the unicast RPF status. If unicast RPF has not been enabled, the **uRPF** flag is not displayed.

On EX3200, EX4200, and EX4300 switches, unicast RPF is implicitly enabled on *all* switch interfaces, including aggregated Ethernet interfaces (also referred to as link aggregation groups or LAGs), integrated routing and bridging (IRB) interfaces, and routed VLAN interfaces (RVIs) when you enable unicast RPF on a single interface. However, the unicast RPF status is shown as enabled only on interfaces for which you have explicitly configured unicast RPF. Thus, the **uRPF** flag is not displayed on interfaces for which you have not explicitly configured unicast RPF even though unicast RPF is implicitly enabled on all interfaces on EX3200 and EX4200 switches.

- Related Documentation**
- [show interfaces xe- on page 368](#)
 - [Example: Configuring Unicast RPF on an EX Series Switch on page 49](#)
 - [Configuring Unicast RPF \(CLI Procedure\) on page 123](#)
 - [Disabling Unicast RPF \(CLI Procedure\) on page 125](#)

- [Troubleshooting Unicast RPF on page 429](#)

Verifying IP Directed Broadcast Status

Purpose	Verify that IP directed broadcast is enabled and is working on the subnet.
Action	Use the show vlans extensive command to verify that IP directed broadcast is enabled and working on the subnet as shown in “Example: Configuring IP Directed Broadcast on a Switch” on page 54 .
Related Documentation	<ul style="list-style-type: none">• Configuring IP Directed Broadcast (CLI Procedure) on page 126• Configuring IP Directed Broadcast (CLI Procedure)• Example: Configuring IP Directed Broadcast on a Switch on page 54

Verifying That Generic Routing Encapsulation Tunneling Is Working Correctly

Purpose	Verify that the generic routing encapsulation (GRE) interface is sending tunneled traffic.
Action	<p>Display status information about the specified GRE interface by using the command show interfaces.</p> <pre>user@switch> show interfaces gr-0/0/0.0 Physical interface: gr-0/0/0, Enabled, Physical link is Up Interface index: 132, SNMP ifIndex: 26 Type: GRE, Link-level type: GRE, MTU: Unlimited, Speed: 800mbps Device flags : Present Running Interface flags: Point-To-Point SNMP-Traps Input rate : 0 bps (0 pps) Output rate : 0 bps (0 pps) Logical interface gr-0/0/0.0 (Index 68) (SNMP ifIndex 47) Flags: Point-To-Point SNMP-Traps 16384 IP-Header 1.1.1.2:1.1.1.1:47:df:64:0000000000000000 Encapsulation: GRE-NULL Input packets : 0 Output packets: 0 Protocol inet, MTU: 1476 Flags: None Addresses, Flags: Is-Primary Local: 1.10.1.1</pre>
Meaning	The output indicates that the GRE interface gr-0/0/0 is up. The output displays the name of the physical interface and the traffic statistics for this interface---the number of and the rate at which input and output bytes and packets are received and transmitted on the physical interface.
Related Documentation	<ul style="list-style-type: none">• Configuring Generic Routing Encapsulation Tunneling (CLI Procedure) on page 130

Verifying That EEE Is Saving Energy on Configured Ports

Purpose	Verify that enabling EEE saves energy on Base-T Copper Ethernet ports.
----------------	--

Action You can see the amount of energy saved by EEE on an EX Series switch using the **show chassis power-budget-statistics** command.

1. View the power budget of an EX Series switch before enabling EEE.

- On an EX6210 switch:

```
user@switch>show chassis power-budget-statistics
PSU 2 (EX6200-PWR-AC2500) : 2500 W Online
PSU 3 ) : 0 W Offline
Total Power supplied by all Online PSUs : 2500 W
Power Redundancy Configuration : N+1
Power Reserved for the Chassis : 500 W

Fan Tray Statistics
FTC 0 : Base power 300 W Power Used nan W

FPC Statistics
power Priority Base power Power Used PoE
FPC 3 (EX6200-48T) : 150 W 61.54 W
0 W 9
FPC 4 (EX6200-SRE64-4XS) : 100 W 48.25 W
0 W 0
FPC 5 (EX6200-SRE64-4XS) : 100 W 48.00 W
0 W 0
FPC 7 (EX6200-48T) : 150 W 63.11 W
0 W 9
FPC 8 (EX6200-48T) : 150 W 12.17 W
0 W 9

Total (non-PoE) Power allocated : 950 W
Total Power allocated for PoE : 0 W
Power Available (Redundant case) : 0 W
Total Power Available : 1550 W
```

- On an EX4300 switch:

```
user@switch>show chassis power-budget-statistics fpc 1
PSU 1 (JPSU-1100-AC-AF0-A) : 1100 W Online
Power redundancy configuration : N+0
Total power supplied by all online PSUs : 1100 W
Base power reserved : 175 W
Non-PoE power being consumed : 95 W
Total Power allocated for PoE : 925 W
Total PoE power consumed : 0 W
Total PoE power remaining : 925 W
```

2. Enable EEE on Base-T Copper Ethernet ports and save the configuration.

3. View the power budget of the switch after enabling EEE.

- On an EX6210 switch:

```
user@switch> show chassis power-budget-statistics
PSU 2 (EX6200-PWR-AC2500) : 2500 W Online
PSU 3 ) : 0 W Offline
Total Power supplied by all Online PSUs : 2500 W
Power Redundancy Configuration : N+1
Power Reserved for the Chassis : 500 W

Fan Tray Statistics
FTC 0 : Base power 300 W Power Used nan W

FPC Statistics
Base power Power Used PoE
```


Source Class Field

For the logical interface, the **Source class** field provides the names of source class usage (SCU) counters per family and per class for a particular interface. The counters display packets and bytes arriving from designated user-selected prefixes. For example:

Source class	Packets (packet-per-second)	Bytes (bits-per-second)
gold	1928095	161959980
(889)	(597762)
bronze	0	0
(0)	(0)
silver	0	0
(0)	(0)

Table 29: request diagnostics tdr Output Fields

Field Name	Field Description
Test Status	<p>Information about the status of the TDR test request:</p> <ul style="list-style-type: none">• Admin Down <i>interface-name</i>—The interface is administratively down. The TDR test cannot run on interfaces that are administratively down.• Interface <i>interface-name</i> not found—The interface does not exist.• Test successfully executed <i>interface-name</i>—The test has successfully started on the interface. You can view the test results with the show diagnostics tdr command.• VCT not supported on <i>interface-name</i>—The TDR test is not supported on the interface.

Sample Output

request diagnostics tdr start interface ge-0/0/19

```
user@switch> request diagnostics tdr start interface ge-0/0/19
```

Interface TDR detail:

Test status : Test successfully executed ge-0/0/19


```
Output:          1694          3          130057          2400
Link:
ge-1/2/0.452
Input :          293          1          26003          1072
Output:         1694          3          130057          2400
ge-1/2/1.452
Input :           0          0           0           0
Output:           0          0           0           0
Marker Statistics:  Marker Rx    Resp Tx    Unknown Rx    Illegal Rx
ge-1/2/0.452             0          0           0           0
ge-1/2/1.452             0          0           0           0
Protocol vpls, MTU: 1518, Generation: 850, Route table: 3
Flags: None
...
```


Table 32: GRE show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
Broadcast	Broadcast address of the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive

Sample Output

show interfaces (GRE)

```

user@host> show interfaces gr-1/2/0
Physical interface: gr-0/0/0, Enabled, Physical link is Up
  Interface index: 132, SNMP ifIndex: 26
  Type: GRE, Link-level type: GRE, MTU: Unlimited, Speed: 800mbps
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)

Logical interface gr-0/0/0.0 (Index 68) (SNMP ifIndex 47)
  Flags: Point-To-Point SNMP-Traps 16384
  IP-Header 1.1.1.2:1.1.1.1:47:df:64:0000000000000000 Encapsulation: GRE-NULL
  Input packets : 0
  Output packets: 0
  Protocol inet, MTU: 1476
  Flags: None
  Addresses, Flags: Is-Primary
    Local: 1.10.1.1

```

show interfaces brief (GRE)

```

user@host> show interfaces gr-1/2/0 brief
Physical interface: gr-1/2/0, Enabled, Physical link is Up
  Type: GRE, Link-level type: GRE, MTU: Unlimited, Speed: 800mbps
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps

Logical interface gr-1/2/0.0
  Flags: Hardware-Down Point-To-Point SNMP-Traps 0x4000
  IP-Header 10.10.0.2:10.10.0.1:47:df:64:0000000000000000
  Encapsulation: GRE-NULL
  inet 10.100.0.1/30
  mpls

```

show interfaces detail (GRE)

```

user@host> show interfaces gr-1/2/0 detail
Physical interface: gr-0/0/0, Enabled, Physical link is Up
  Interface index: 132, SNMP ifIndex: 26, Generation: 13
  Type: GRE, Link-level type: GRE, MTU: Unlimited, Speed: 800mbps
  Hold-times      : Up 0 ms, Down 0 ms
  Device flags    : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Statistics last cleared: Never
  Traffic statistics:
    Input bytes   : 0 0 bps
    Output bytes  : 0 0 bps
    Input packets : 0 0 pps
    Output packets: 0 0 pps

Logical interface gr-0/0/0.0 (Index 68) (SNMP ifIndex 47) (Generation 8)
  Flags: Point-To-Point SNMP-Traps 16384
  IP-Header 1.1.1.2:1.1.1.1:47:df:64:0000000000000000 Encapsulation: GRE-NULL
  Traffic statistics:
    Input bytes   : 0
    Output bytes  : 0
    Input packets : 0

```



```
Copy-tos-to-outer-ip-header: Off
Gre keepalives configured: Off, Gre keepalives adjacency state: down
Traffic statistics:
  Input bytes :          260
  Output bytes :       2880148
  Input packets:          4
  Output packets:      10002
Local statistics:
  Input bytes :          112
  Output bytes :           0
  Input packets:         2
  Output packets:        0
Transit statistics:
  Input bytes :          148          0 bps
  Output bytes :       2880148          0 bps
  Input packets:         2          0 pps
  Output packets:      10002          0 pps
Protocol inet, Generation: 171, Route table: 0
Flags: None
Addresses, Flags: Is-Preferred Is-Primary
  Destination: 70.70.70/24, Local: 70.70.70.10, Broadcast: 70.70.70.255,
  Generation: 160
```

[show interfaces extensive \(GRE\)](#)

The output for the **show interfaces extensive** command is identical to that for the **show interfaces detail** command. For sample output, see [show interfaces detail \(GRE\) on page 331](#) and [show interfaces detail \(GRE\) on an EX4200 Virtual Chassis Member Switch on page 333](#).

Table 33: show interfaces diagnostics optics Output Fields (*continued*)

Field Name	Field Description
Laser output power low warning (Not available for QSFP+ transceivers)	Displays whether the laser output power low warning is On or Off .
Laser temperature high alarm (Not available for SFP, SFP+, XFP, and QSFP+ transceivers)	Displays whether the laser temperature high alarm is On or Off .
Laser temperature low alarm (Not available for SFP, SFP+, XFP, and QSFP+ transceivers)	Displays whether the laser temperature low alarm is On or Off .
Laser temperature high warning (Not available for SFP, SFP+, XFP, and QSFP+ transceivers)	Displays whether the laser temperature high warning is On or Off .
Laser temperature low warning (Not available for SFP, SFP+, XFP, and QSFP+ transceivers)	Displays whether the laser temperature low warning is On or Off .
Module temperature high alarm (Not available for QSFP+ transceivers)	Displays whether the module temperature high alarm is On or Off .
Module temperature low alarm (Not available for QSFP+ transceivers)	Displays whether the module temperature low alarm is On or Off .
Module temperature high warning (Not available for QSFP+ transceivers)	Displays whether the module temperature high warning is On or Off .
Module temperature low warning (Not available for QSFP+ transceivers)	Displays whether the module temperature low warning is On or Off .
Module voltage high alarm (Not available for XFP and QSFP+ transceivers)	Displays whether the module voltage high alarm is On or Off .
Module voltage low alarm (Not available for XFP and QSFP+ transceivers)	Displays whether the module voltage low alarm is On or Off .
Module voltage high warning (Not available for XFP and QSFP+ transceivers)	Displays whether the module voltage high warning is On or Off .
Module voltage low warning (Not available for XFP and QSFP+ transceivers)	Displays whether the module voltage low warning is On or Off .

Sample Output

show interfaces diagnostics optics et-3/0/0 (QSFP+ Transceiver)

```

user@switch> show interfaces diagnostics optics et-3/0/0
Physical interface: et-3/0/0
  Module temperature           : 33 degrees C / 92 degrees F
  Module voltage               : 3.3060 V
Lane 0
  Laser bias current           : 7.182 mA
  Laser receiver power         : 0.743 mW / -1.29 dBm
  Laser bias current high alarm : Off
  Laser bias current low alarm  : Off
  Laser bias current high warning : Off
  Laser bias current low warning : Off
  Laser receiver power high alarm : Off
  Laser receiver power low alarm  : Off
  Laser receiver power high warning : Off
  Laser receiver power low warning : Off
  Tx loss of signal functionality alarm : Off
  Rx loss of signal alarm        : Off
Lane 1
  Laser bias current           : 7.326 mA
  Laser receiver power         : 0.752 mW / -1.24 dBm
  Laser bias current high alarm : Off
  Laser bias current low alarm  : Off
  Laser bias current high warning : Off
  Laser bias current low warning : Off
  Laser receiver power high alarm : Off
  Laser receiver power low alarm  : Off
  Laser receiver power high warning : Off
  Laser receiver power low warning : Off
  Tx loss of signal functionality alarm : Off
  Rx loss of signal alarm        : Off
Lane 2
  Laser bias current           : 7.447 mA
  Laser receiver power         : 0.790 mW / -1.03 dBm
  Laser bias current high alarm : Off
  Laser bias current low alarm  : Off
  Laser bias current high warning : Off
  Laser bias current low warning : Off
  Laser receiver power high alarm : Off
  Laser receiver power low alarm  : Off
  Laser receiver power high warning : Off
  Laser receiver power low warning : Off
  Tx loss of signal functionality alarm : Off
  Rx loss of signal alarm        : Off
Lane 3
  Laser bias current           : 7.734 mA
  Laser receiver power         : 0.768 mW / -1.15 dBm
  Laser bias current high alarm : Off
  Laser bias current low alarm  : Off
  Laser bias current high warning : Off
  Laser bias current low warning : Off
  Laser receiver power high alarm : Off
  Laser receiver power low alarm  : Off
  Laser receiver power high warning : Off
  Laser receiver power low warning : Off
  Tx loss of signal functionality alarm : Off
  Rx loss of signal alarm        : Off

```



```
Input  packets:           0
Output packets:           0
Protocol eth-switch, Generation: 159, Route table: 0
Flags: None
Input Filters: f2,
Output Filters: f1,,,,
```



```
Flags: SNMP-Traps 0x0 Encapsulation: ENET2
Traffic statistics:
  Input bytes :          0
  Output bytes :        9375416
  Input packets:          0
  Output packets:       48901
Local statistics:
  Input bytes :          0
  Output bytes :        9375416
  Input packets:          0
  Output packets:       48901
Transit statistics:
  Input bytes :          0          0 bps
  Output bytes :          0          0 bps
  Input packets:          0          0 pps
  Output packets:          0          0 pps
Protocol eth-switch, Generation: 1937, Route table: 0
  Flags: Trunk-Mode
```



```

Bytes : 0
Tail-dropped packets : 0

```

show interfaces queue xe-6/0/39 (Line Card with Oversubscribed Ports in an EX8200 Switch)

```

user@switch> show interfaces queue xe-6/0/39

Physical interface: xe-6/0/39, Enabled, Physical link is Up
  Interface index: 291, SNMP ifIndex: 1641
Forwarding classes: 16 supported, 7 in use
Ingress queues: 1 supported, 1 in use
  Transmitted:
    Packets : 337069086018
    Bytes : 43144843010304
    Tail-dropped packets : 8003867575
PFE chassis queues: 1 supported, 1 in use
  Transmitted:
    Packets : 0
    Bytes : 0
    Tail-dropped packets : 0
Forwarding classes: 16 supported, 7 in use
Egress queues: 8 supported, 7 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
  Transmitted:
    Packets : 334481399932
    Bytes : 44151544791024
    Tail-dropped packets : 0
Queue: 1, Forwarding classes: assured-forwarding
  Queued:
  Transmitted:
    Packets : 0
    Bytes : 0
    Tail-dropped packets : 0
Queue: 2, Forwarding classes: mcast-be
  Queued:
  Transmitted:
    Packets : 274948977
    Bytes : 36293264964
    Tail-dropped packets : 0
Queue: 4, Forwarding classes: mcast-ef
  Queued:
  Transmitted:
    Packets : 0
    Bytes : 0
    Tail-dropped packets : 0
Queue: 5, Forwarding classes: expedited-forwarding
  Queued:
  Transmitted:
    Packets : 0
    Bytes : 0
    Tail-dropped packets : 0
Queue: 6, Forwarding classes: mcast-af
  Queued:
  Transmitted:
    Packets : 0
    Bytes : 0
    Tail-dropped packets : 0
Queue: 7, Forwarding classes: network-control
  Queued:
  Transmitted:

```


Table 38: show interfaces vlan Output Fields (*continued*)

Field Name	Field Description	Level of Output
Input Errors	<p>Input errors on the interface. The following paragraphs explain some of the counters whose meaning may not be obvious.</p> <ul style="list-style-type: none"> • Errors—Sum of the incoming frame aborts and FCS errors. • Drops—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this value increments once for every packet that is dropped by the ASIC's RED mechanism. • Framing errors—Number of packets received with an invalid frame checksum (FCS). • Runts—Number of frames received that are smaller than the runt threshold. • Policed discards—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle. • L3 incompletes—Number of incoming packets discarded because they failed Layer 3 sanity checks of the headers. For example, a frame with less than 20 bytes of available IP header is discarded. • L2 channel errors—Number of times the software did not find a valid logical interface for an incoming frame. • L2 mismatch timeouts—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • FIFO errors—Number of FIFO errors in the receive direction that are reported by the ASIC. If this value is ever nonzero, the interface is probably malfunctioning. • Resource errors—Sum of transmit drops. 	extensive


```
Transit statistics:
Input bytes :          0          0 bps
Output bytes :          0          0 bps
Input packets:         0          0 pps
Output packets:        0          0 pps
Protocol inet, Generation: 159, Route table: 0
  Flags: Targeted-broadcast
  Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
    Destination: 10.1.1/24, Local: 10.1.1.1, Broadcast: 10.1.1.255,
    Generation: 138

Logical interface vlan.1 (Index 83) (SNMP ifIndex 558) (Generation 148)
  Flags: Link-Layer-Down SNMP-Traps 0x0 Encapsulation: ENET2
  Traffic statistics:
    Input bytes :          0
    Output bytes :         42
    Input packets:         0
    Output packets:        1
  Local statistics:
    Input bytes :          0
    Output bytes :         42
    Input packets:         0
    Output packets:        1
  Transit statistics:
    Input bytes :          0          0 bps
    Output bytes :          0          0 bps
    Input packets:         0          0 pps
    Output packets:        0          0 pps
  Protocol inet, Generation: 160, Route table: 0
    Flags: Targeted-broadcast
    Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
      Destination: 10.1.2/24, Local: 10.1.2.1, Broadcast: 10.1.2.255,
      Generation: 140
```


xe-4/0/7	Actor	No	No	No	No	No	Yes	Fast	Active
xe-4/0/7	Partner	No	No	No	Yes	Yes	Yes	Fast	Passive

LACP protocol:	Receive State	Transmit State	Mux State
xe-2/0/7(Active)	Current	Fast periodic	Collecting distributing
xe-34/0/7(Standby)	Current	Fast periodic	Waiting

show lacp interfaces (QFX Series)

```
user@switch> show lacp interfaces nodegroup1:ae0 extensive
```

```
Aggregated interface: nodegroup1:ae0
```

LACP state:	Role	Exp	Def	Dist	Col	Syn	Aggr	Timeout	Activity
node1:xe-0/0/1FUP	Actor		No	Yes	No	No	No	Yes	Fast
Active									
node1xe-0/0/1FUP	Partner		No	Yes	No	No	No	Yes	Fast
Passive									
node2:xe-0/0/2	Actor		No	Yes	No	No	No	Yes	Fast
Active									
node2:xe-0/0/2	Partner		No	Yes	No	No	No	Yes	Fast
Passive									

	LACP protocol:	Receive State	Transmit State	Mux State
	node1:xe-0/0/1FUP	Current	Fast periodic	Collecting
distributing	node2:xe-0/0/2	Current	Fast periodic	Collecting
distributing	node1:xe-0/0/1 (active)	Current	Fast periodic	Collecting
distributing	node2:xe-0/0/2 (standby)	Current	Fast periodic	WAITING

show virtual-chassis vc-port diagnostics optics

Syntax	<pre>show virtual-chassis vc-port diagnostics optics <all-members> <interface-name> <local> <member member-id></pre>
Release Information	<p>Command introduced in Junos OS Release 12.2 for EX Series switches.</p> <p>Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).</p>
Description	<p>Display diagnostics data and alarms for Ethernet optical transceivers installed in ports configured as Virtual Chassis Ports (VCPs) in an EX Series switches. The information provided by this command is known as digital optical monitoring (DOM) information.</p> <p>Thresholds that trigger a high alarm, low alarm, high warning, or low warning are set by the transponder vendors. Generally, a high alarm or low alarm indicates that a transceiver is not operating properly. DOM information can be used to diagnose why a transceiver is not working.</p> <p>On some EX Series switches, the request virtual-chassis vc-port diagnostics optics command must be entered to run a diagnostic scan before you can gather the show virtual-chassis vc-port diagnostics optics output.</p>
Options	<p>none—Display diagnostics information for transceivers installed in VCPs of all members of a Virtual Chassis or VCF.</p> <p>all-members—(Optional) Display diagnostics information for transceivers installed in VCPs of all members of a Virtual Chassis or VCF.</p> <p>interface-name—(Optional) Display diagnostics information for the transceiver installed in a specified VCP.</p> <p>local—(Optional) Display diagnostics information for transceivers installed in VCPs on the switch or external Routing Engine on which this command is entered.</p> <p>member member-id—(Optional) Display diagnostics information for transceivers installed in VCPs on a specified member of a Virtual Chassis or VCF.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • <i>show virtual-chassis vc-port</i> • <i>Installing a Transceiver in a Switch</i> • <i>Removing a Transceiver from a Switch</i> • Junos OS Ethernet Interfaces Configuration Guide
List of Sample Output	<p>show virtual-chassis vc-port diagnostics optics on page 408</p> <p>show virtual-chassis vc-port diagnostics optics (interface-name) on page 413</p>

[show virtual-chassis vc-port diagnostics optics local on page 415](#)

[show virtual-chassis vc-port diagnostics optics \(member member-id\) on page 417](#)

Output Fields [Table 40 on page 406](#) lists the output fields for the **show virtual-chassis vc-port diagnostics optics** command. Output fields are listed in the approximate order in which they appear.

Table 40: show virtual-chassis vc-port diagnostics optics Output Fields

Field Name	Field Description
FPC	Displays the FPC slot number.
Virtual chassis port	Displays the name of the VCP.
Laser bias current	Displays the magnitude of the laser bias power setting current, in milliamperes (mA). The laser bias provides direct modulation of laser diodes and modulates currents.
Laser output power	Displays the laser output power, in milliwatts (mW) and decibels referred to 1.0 mW (dBm).
Module temperature	Displays the temperature, in Celsius and Fahrenheit.
Module voltage	Displays the voltage, in Volts.
Receiver signal average optical power	Displays the receiver signal average optical power, in milliwatts (mW) and decibels referred to 1.0 mW (dBm).
Laser bias current high alarm	Displays whether the laser bias power setting high alarm is <i>On</i> or <i>Off</i> .
Laser bias current low alarm	Displays whether the laser bias power setting low alarm is <i>On</i> or <i>Off</i> .
Laser bias current high warning	Displays whether the laser bias power setting high warning is <i>On</i> or <i>Off</i> .
Laser bias current low warning	Displays whether the laser bias power setting low warning is <i>On</i> or <i>Off</i> .
Laser output power high alarm	Displays whether the laser output power high alarm is <i>On</i> or <i>Off</i> .
Laser output power low alarm	Displays whether the laser output power low alarm is <i>On</i> or <i>Off</i> .
Laser output power high warning	Displays whether the laser output power high warning is <i>On</i> or <i>Off</i> .
Laser output power low warning	Displays whether the laser output power low warning is <i>On</i> or <i>Off</i> .
Module temperature high alarm	Displays whether the module temperature high alarm is <i>On</i> or <i>Off</i> .
Module temperature low alarm	Displays whether the module temperature low alarm is <i>On</i> or <i>Off</i> .
Module temperature high warning	Displays whether the module temperature high warning is <i>On</i> or <i>Off</i> .
Module temperature low warning	Displays whether the module temperature low warning is <i>On</i> or <i>Off</i> .

Table 40: show virtual-chassis vc-port diagnostics optics Output Fields (*continued*)

Field Name	Field Description
Module voltage high alarm	Displays whether the module voltage high alarm is <i>On</i> or <i>Off</i> .
Module voltage low alarm	Displays whether the module voltage low alarm is <i>On</i> or <i>Off</i> .
Module voltage high warning	Displays whether the module voltage high warning is <i>On</i> or <i>Off</i> .
Module voltage low warning	Displays whether the module voltage low warning is <i>On</i> or <i>Off</i> .
Laser rx power high alarm	Displays whether the receive laser power high alarm is <i>On</i> or <i>Off</i> .
Laser rx power low alarm	Displays whether the receive laser power low alarm is <i>On</i> or <i>Off</i> .
Laser rx power high warning	Displays whether the receive laser power high warning is <i>On</i> or <i>Off</i> .
Laser rx power low warning	Displays whether the receive laser power low warning is <i>On</i> or <i>Off</i> .
Laser bias current high alarm threshold	Displays the vendor-specified threshold for the laser bias current high alarm.
Laser bias current low alarm threshold	Displays the vendor-specified threshold for the laser bias current low alarm.
Laser bias current high warning threshold	Displays the vendor-specified threshold for the laser bias current high warning.
Laser bias current low warning threshold	Displays the vendor-specified threshold for the laser bias current low warning.
Laser output power high alarm threshold	Displays the vendor-specified threshold for the laser output power high alarm.
Laser output power low alarm threshold	Displays the vendor-specified threshold for the laser output power low alarm.
Laser output power high warning threshold	Displays the vendor-specified threshold for the laser output power high warning.
Laser output power low warning threshold	Displays the vendor-specified threshold for the laser output power low warning.
Module temperature high alarm threshold	Displays the vendor-specified threshold for the module temperature high alarm.
Module temperature low alarm threshold	Displays the vendor-specified threshold for the module temperature low alarm.
Module temperature high warning threshold	Displays the vendor-specified threshold for the module temperature high warning.
Module temperature low warning threshold	Displays the vendor-specified threshold for the module temperature low warning.

Table 40: show virtual-chassis vc-port diagnostics optics Output Fields (*continued*)

Field Name	Field Description
Module voltage high alarm threshold	Displays the vendor-specified threshold for the module voltage high alarm.
Module voltage low alarm threshold	Displays the vendor-specified threshold for the module voltage low alarm.
Module voltage high warning threshold	Displays the vendor-specified threshold for the module voltage high warning.
Module voltage low warning threshold	Displays the vendor-specified threshold for the module voltage low warning.
Laser rx power high alarm threshold	Displays the vendor-specified threshold for the laser rx power high alarm.
Laser rx power low alarm threshold	Displays the vendor-specified threshold for the laser rx power low alarm.
Laser rx power high warning threshold	Displays the vendor-specified threshold for the laser rx power high warning.
Laser rx power low warning threshold	Displays the vendor-specified threshold for the laser rx power low warning.

Sample Output

show virtual-chassis vc-port diagnostics optics

```

user@switch> show virtual-chassis vc-port diagnostics optics
fpc0:
-----
Virtual chassis port: vcp-0
  Optical diagnostics                : N/A
Virtual chassis port: vcp-1
  Optical diagnostics                : N/A

fpc1:
-----
Virtual chassis port: vcp-0
  Optical diagnostics                : N/A
Virtual chassis port: vcp-1
  Optical diagnostics                : N/A

fpc2:
-----
Virtual chassis port: vcp-2/0
  Optical diagnostics                : N/A
Virtual chassis port: vcp-2/1
  Optical diagnostics                : N/A
Virtual chassis port: vcp-255/0/14
  Optical diagnostics                : N/A
Virtual chassis port: vcp-255/0/15
  Optical diagnostics                : N/A
Virtual chassis port: vcp-255/0/24
  Laser bias current                 : 4.130 mA
  Laser output power                 : 0.2450 mW / -6.11 dBm
  Module temperature                 : 32 degrees C / 90 degrees F
  Module voltage                     : 3.3530 V
  Receiver signal average optical power : 0.0971 mW / -10.13 dBm
  Laser bias current high alarm      : Off
  Laser bias current low alarm       : Off

```

```

Laser bias current high warning      : Off
Laser bias current low warning       : Off
Laser output power high alarm        : Off
Laser output power low alarm         : Off
Laser output power high warning      : Off
Laser output power low warning       : Off
Module temperature high alarm        : Off
Module temperature low alarm         : Off
Module temperature high warning      : Off
Module temperature low warning       : Off
Module voltage high alarm            : Off
Module voltage low alarm             : Off
Module voltage high warning          : Off
Module voltage low warning           : Off
Laser rx power high alarm            : Off
Laser rx power low alarm             : Off
Laser rx power high warning          : Off
Laser rx power low warning           : Off
Laser bias current high alarm threshold : 14.998 mA
Laser bias current low alarm threshold : 0.998 mA
Laser bias current high warning threshold : 14.000 mA
Laser bias current low warning threshold : 1.198 mA
Laser output power high alarm threshold : 0.7940 mW / -1.00 dBm
Laser output power low alarm threshold : 0.0790 mW / -11.02 dBm
Laser output power high warning threshold : 0.6300 mW / -2.01 dBm
Laser output power low warning threshold : 0.0990 mW / -10.04 dBm
Module temperature high alarm threshold : 85 degrees C / 185 degrees F
Module temperature low alarm threshold : -10 degrees C / 14 degrees F
Module temperature high warning threshold : 80 degrees C / 176 degrees F
Module temperature low warning threshold : -5 degrees C / 23 degrees F
Module voltage high alarm threshold : 3.600 V
Module voltage low alarm threshold : 3.000 V
Module voltage high warning threshold : 3.499 V
Module voltage low warning threshold : 3.099 V
Laser rx power high alarm threshold : 1.5848 mW / 2.00 dBm
Laser rx power low alarm threshold : 0.0100 mW / -20.00 dBm
Laser rx power high warning threshold : 1.2589 mW / 1.00 dBm
Laser rx power low warning threshold : 0.0125 mW / -19.03 dBm
Virtual chassis port: vcp-255/0/3
Laser bias current                  : 5.428 mA
Laser output power                  : 0.4760 mW / -3.22 dBm
Module temperature                  : 28 degrees C / 83 degrees F
Module voltage                      : 3.3440 V
Receiver signal average optical power : 0.4002 mW / -3.98 dBm
Laser bias current high alarm       : Off
Laser bias current low alarm        : Off
Laser bias current high warning     : Off
Laser bias current low warning      : Off
Laser output power high alarm       : Off
Laser output power low alarm        : Off
Laser output power high warning     : Off
Laser output power low warning      : Off
Module temperature high alarm       : Off
Module temperature low alarm        : Off
Module temperature high warning     : Off
Module temperature low warning      : Off
Module voltage high alarm           : Off
Module voltage low alarm            : Off
Module voltage high warning         : Off
Module voltage low warning          : Off
Laser rx power high alarm           : Off

```

```

Laser rx power low alarm           : Off
Laser rx power high warning        : Off
Laser rx power low warning         : Off
Laser bias current high alarm threshold : 10.500 mA
Laser bias current low alarm threshold : 2.000 mA
Laser bias current high warning threshold : 9.000 mA
Laser bias current low warning threshold : 2.500 mA
Laser output power high alarm threshold : 1.4120 mW / 1.50 dBm
Laser output power low alarm threshold : 0.0740 mW / -11.31 dBm
Laser output power high warning threshold : 0.7070 mW / -1.51 dBm
Laser output power low warning threshold : 0.1860 mW / -7.30 dBm
Module temperature high alarm threshold : 75 degrees C / 167 degrees F
Module temperature low alarm threshold : -5 degrees C / 23 degrees F
Module temperature high warning threshold : 70 degrees C / 158 degrees F
Module temperature low warning threshold : 0 degrees C / 32 degrees F
Module voltage high alarm threshold : 3.630 V
Module voltage low alarm threshold : 2.970 V
Module voltage high warning threshold : 3.465 V
Module voltage low warning threshold : 3.135 V
Laser rx power high alarm threshold : 1.5849 mW / 2.00 dBm
Laser rx power low alarm threshold : 0.0407 mW / -13.90 dBm
Laser rx power high warning threshold : 0.7943 mW / -1.00 dBm
Laser rx power low warning threshold : 0.1023 mW / -9.90 dBm

```

fpc3:

Virtual chassis port: vcp-255/0/2

```

Laser bias current           : 7.876 mA
Laser output power           : 0.5330 mW / -2.73 dBm
Module temperature            : 26 degrees C / 78 degrees F
Module voltage                : 3.3060 V
Receiver signal average optical power : 0.4885 mW / -3.11 dBm
Laser bias current high alarm : Off
Laser bias current low alarm  : Off
Laser bias current high warning : Off
Laser bias current low warning : Off
Laser output power high alarm  : Off
Laser output power low alarm   : Off
Laser output power high warning : Off
Laser output power low warning : Off
Module temperature high alarm  : Off
Module temperature low alarm   : Off
Module temperature high warning : Off
Module temperature low warning : Off
Module voltage high alarm      : Off
Module voltage low alarm       : Off
Module voltage high warning    : Off
Module voltage low warning     : Off
Laser rx power high alarm      : Off
Laser rx power low alarm       : Off
Laser rx power high warning    : Off
Laser rx power low warning     : Off
Laser bias current high alarm threshold : 14.500 mA
Laser bias current low alarm threshold : 3.500 mA
Laser bias current high warning threshold : 14.500 mA
Laser bias current low warning threshold : 3.500 mA
Laser output power high alarm threshold : 1.8620 mW / 2.70 dBm
Laser output power low alarm threshold : 0.0740 mW / -11.31 dBm
Laser output power high warning threshold : 0.7410 mW / -1.30 dBm
Laser output power low warning threshold : 0.1860 mW / -7.30 dBm
Module temperature high alarm threshold : 75 degrees C / 167 degrees F

```

```

Module temperature low alarm threshold : -5 degrees C / 23 degrees F
Module temperature high warning threshold : 70 degrees C / 158 degrees F
Module temperature low warning threshold : 0 degrees C / 32 degrees F
Module voltage high alarm threshold : 3.630 V
Module voltage low alarm threshold : 2.970 V
Module voltage high warning threshold : 3.465 V
Module voltage low warning threshold : 3.135 V
Laser rx power high alarm threshold : 1.9952 mW / 3.00 dBm
Laser rx power low alarm threshold : 0.0407 mW / -13.90 dBm
Laser rx power high warning threshold : 0.7943 mW / -1.00 dBm
Laser rx power low warning threshold : 0.1023 mW / -9.90 dBm
Virtual chassis port: vcp-255/0/3
Laser bias current : 5.052 mA
Laser output power : 0.5030 mW / -2.98 dBm
Module temperature : 24 degrees C / 75 degrees F
Module voltage : 3.2890 V
Receiver signal average optical power : 0.5028 mW / -2.99 dBm
Laser bias current high alarm : Off
Laser bias current low alarm : Off
Laser bias current high warning : Off
Laser bias current low warning : Off
Laser output power high alarm : Off
Laser output power low alarm : Off
Laser output power high warning : Off
Laser output power low warning : Off
Module temperature high alarm : Off
Module temperature low alarm : Off
Module temperature high warning : Off
Module temperature low warning : Off
Module voltage high alarm : Off
Module voltage low alarm : Off
Module voltage high warning : Off
Module voltage low warning : Off
Laser rx power high alarm : Off
Laser rx power low alarm : Off
Laser rx power high warning : Off
Laser rx power low warning : Off
Laser bias current high alarm threshold : 10.500 mA
Laser bias current low alarm threshold : 2.000 mA
Laser bias current high warning threshold : 9.000 mA
Laser bias current low warning threshold : 2.500 mA
Laser output power high alarm threshold : 1.4120 mW / 1.50 dBm
Laser output power low alarm threshold : 0.0740 mW / -11.31 dBm
Laser output power high warning threshold : 0.7070 mW / -1.51 dBm
Laser output power low warning threshold : 0.1860 mW / -7.30 dBm
Module temperature high alarm threshold : 75 degrees C / 167 degrees F
Module temperature low alarm threshold : -5 degrees C / 23 degrees F
Module temperature high warning threshold : 70 degrees C / 158 degrees F
Module temperature low warning threshold : 0 degrees C / 32 degrees F
Module voltage high alarm threshold : 3.630 V
Module voltage low alarm threshold : 2.970 V
Module voltage high warning threshold : 3.465 V
Module voltage low warning threshold : 3.135 V
Laser rx power high alarm threshold : 1.5849 mW / 2.00 dBm
Laser rx power low alarm threshold : 0.0407 mW / -13.90 dBm
Laser rx power high warning threshold : 0.7943 mW / -1.00 dBm
Laser rx power low warning threshold : 0.1023 mW / -9.90 dBm
Virtual chassis port: vcp-255/0/4
Laser bias current : 7.978 mA
Laser output power : 0.5460 mW / -2.63 dBm
Module temperature : 24 degrees C / 76 degrees F

```

```

Module voltage                                     : 3.3060 V
Receiver signal average optical power             : 0.6305 mW / -2.00 dBm
Laser bias current high alarm                     : Off
Laser bias current low alarm                      : Off
Laser bias current high warning                   : Off
Laser bias current low warning                    : Off
Laser output power high alarm                     : Off
Laser output power low alarm                      : Off
Laser output power high warning                   : Off
Laser output power low warning                    : Off
Module temperature high alarm                     : Off
Module temperature low alarm                      : Off
Module temperature high warning                   : Off
Module temperature low warning                    : Off
Module voltage high alarm                         : Off
Module voltage low alarm                          : Off
Module voltage high warning                       : Off
Module voltage low warning                        : Off
Laser rx power high alarm                         : Off
Laser rx power low alarm                          : Off
Laser rx power high warning                       : Off
Laser rx power low warning                        : Off
Laser bias current high alarm threshold           : 14.500 mA
Laser bias current low alarm threshold             : 3.500 mA
Laser bias current high warning threshold         : 14.500 mA
Laser bias current low warning threshold          : 3.500 mA
Laser output power high alarm threshold           : 1.8620 mW / 2.70 dBm
Laser output power low alarm threshold            : 0.0740 mW / -11.31 dBm
Laser output power high warning threshold         : 0.7410 mW / -1.30 dBm
Laser output power low warning threshold          : 0.1860 mW / -7.30 dBm
Module temperature high alarm threshold           : 75 degrees C / 167 degrees F
Module temperature low alarm threshold             : -5 degrees C / 23 degrees F
Module temperature high warning threshold         : 70 degrees C / 158 degrees F
Module temperature low warning threshold          : 0 degrees C / 32 degrees F
Module voltage high alarm threshold               : 3.630 V
Module voltage low alarm threshold                : 2.970 V
Module voltage high warning threshold             : 3.465 V
Module voltage low warning threshold              : 3.135 V
Laser rx power high alarm threshold               : 1.9952 mW / 3.00 dBm
Laser rx power low alarm threshold                : 0.0407 mW / -13.90 dBm
Laser rx power high warning threshold             : 0.7943 mW / -1.00 dBm
Laser rx power low warning threshold              : 0.1023 mW / -9.90 dBm

```

fpc4:

```

-----
Virtual chassis port: vcp-0
  Optical diagnostics                               : N/A
Virtual chassis port: vcp-1
  Optical diagnostics                               : N/A
Virtual chassis port: vcp-255/0/4
  Laser bias current                               : 7.860 mA
  Laser output power                               : 0.5370 mW / -2.70 dBm
  Module temperature                               : 24 degrees C / 75 degrees F
  Module voltage                                   : 3.2920 V
  Receiver signal average optical power            : 0.6271 mW / -2.03 dBm
  Laser bias current high alarm                    : Off
  Laser bias current low alarm                     : Off
  Laser bias current high warning                  : Off
  Laser bias current low warning                   : Off
  Laser output power high alarm                    : Off
  Laser output power low alarm                     : Off
  Laser output power high warning                  : Off

```

```

Laser output power low warning      : Off
Module temperature high alarm       : Off
Module temperature low alarm        : Off
Module temperature high warning     : Off
Module temperature low warning      : Off
Module voltage high alarm           : Off
Module voltage low alarm            : Off
Module voltage high warning         : Off
Module voltage low warning          : Off
Laser rx power high alarm           : Off
Laser rx power low alarm            : Off
Laser rx power high warning         : Off
Laser rx power low warning          : Off
Laser bias current high alarm threshold : 14.500 mA
Laser bias current low alarm threshold : 3.500 mA
Laser bias current high warning threshold : 14.500 mA
Laser bias current low warning threshold : 3.500 mA
Laser output power high alarm threshold : 1.8620 mW / 2.70 dBm
Laser output power low alarm threshold : 0.0740 mW / -11.31 dBm
Laser output power high warning threshold : 0.7410 mW / -1.30 dBm
Laser output power low warning threshold : 0.1860 mW / -7.30 dBm
Module temperature high alarm threshold : 75 degrees C / 167 degrees F
Module temperature low alarm threshold : -5 degrees C / 23 degrees F
Module temperature high warning threshold : 70 degrees C / 158 degrees F
Module temperature low warning threshold : 0 degrees C / 32 degrees F
Module voltage high alarm threshold : 3.630 V
Module voltage low alarm threshold : 2.970 V
Module voltage high warning threshold : 3.465 V
Module voltage low warning threshold : 3.135 V
Laser rx power high alarm threshold : 1.9952 mW / 3.00 dBm
Laser rx power low alarm threshold : 0.0407 mW / -13.90 dBm
Laser rx power high warning threshold : 0.7943 mW / -1.00 dBm
Laser rx power low warning threshold : 0.1023 mW / -9.90 dBm

```

show virtual-chassis vc-port diagnostics optics (interface-name)

```

user@external-routing-engine> show virtual-chassis vc-port diagnostics optics vcp-255/0/3
fpc0:
-----

```

```

fpc1:
-----

```

```

fpc2:
-----

```

```

Virtual chassis port: vcp-255/0/3
Laser bias current      : 5.448 mA
Laser output power      : 0.4770 mW / -3.21 dBm
Module temperature      : 28 degrees C / 82 degrees F
Module voltage          : 3.3450 V
Receiver signal average optical power : 0.3973 mW / -4.01 dBm
Laser bias current high alarm : Off
Laser bias current low alarm  : Off
Laser bias current high warning : Off
Laser bias current low warning : Off
Laser output power high alarm : Off
Laser output power low alarm  : Off
Laser output power high warning : Off
Laser output power low warning : Off
Module temperature high alarm : Off
Module temperature low alarm  : Off

```

```

Module temperature high warning      : Off
Module temperature low warning      : Off
Module voltage high alarm           : Off
Module voltage low alarm            : Off
Module voltage high warning         : Off
Module voltage low warning          : Off
Laser rx power high alarm           : Off
Laser rx power low alarm            : Off
Laser rx power high warning         : Off
Laser rx power low warning          : Off
Laser bias current high alarm threshold : 10.500 mA
Laser bias current low alarm threshold : 2.000 mA
Laser bias current high warning threshold : 9.000 mA
Laser bias current low warning threshold : 2.500 mA
Laser output power high alarm threshold : 1.4120 mW / 1.50 dBm
Laser output power low alarm threshold : 0.0740 mW / -11.31 dBm
Laser output power high warning threshold : 0.7070 mW / -1.51 dBm
Laser output power low warning threshold : 0.1860 mW / -7.30 dBm
Module temperature high alarm threshold : 75 degrees C / 167 degrees F
Module temperature low alarm threshold : -5 degrees C / 23 degrees F
Module temperature high warning threshold : 70 degrees C / 158 degrees F
Module temperature low warning threshold : 0 degrees C / 32 degrees F
Module voltage high alarm threshold : 3.630 V
Module voltage low alarm threshold : 2.970 V
Module voltage high warning threshold : 3.465 V
Module voltage low warning threshold : 3.135 V
Laser rx power high alarm threshold : 1.5849 mW / 2.00 dBm
Laser rx power low alarm threshold : 0.0407 mW / -13.90 dBm
Laser rx power high warning threshold : 0.7943 mW / -1.00 dBm
Laser rx power low warning threshold : 0.1023 mW / -9.90 dBm

```

fpc3:

Virtual chassis port: vcp-255/0/3

```

Laser bias current                  : 5.040 mA
Laser output power                  : 0.5020 mW / -2.99 dBm
Module temperature                  : 24 degrees C / 74 degrees F
Module voltage                      : 3.2870 V
Receiver signal average optical power : 0.5073 mW / -2.95 dBm
Laser bias current high alarm       : Off
Laser bias current low alarm        : Off
Laser bias current high warning     : Off
Laser bias current low warning      : Off
Laser output power high alarm       : Off
Laser output power low alarm        : Off
Laser output power high warning     : Off
Laser output power low warning      : Off
Module temperature high alarm       : Off
Module temperature low alarm        : Off
Module temperature high warning     : Off
Module temperature low warning      : Off
Module voltage high alarm           : Off
Module voltage low alarm            : Off
Module voltage high warning         : Off
Module voltage low warning          : Off
Laser rx power high alarm           : Off
Laser rx power low alarm            : Off
Laser rx power high warning         : Off
Laser rx power low warning          : Off
Laser bias current high alarm threshold : 10.500 mA
Laser bias current low alarm threshold : 2.000 mA

```

```

Laser bias current high warning threshold : 9.000 mA
Laser bias current low warning threshold  : 2.500 mA
Laser output power high alarm threshold   : 1.4120 mW / 1.50 dBm
Laser output power low alarm threshold    : 0.0740 mW / -11.31 dBm
Laser output power high warning threshold : 0.7070 mW / -1.51 dBm
Laser output power low warning threshold  : 0.1860 mW / -7.30 dBm
Module temperature high alarm threshold   : 75 degrees C / 167 degrees F
Module temperature low alarm threshold    : -5 degrees C / 23 degrees F
Module temperature high warning threshold : 70 degrees C / 158 degrees F
Module temperature low warning threshold  : 0 degrees C / 32 degrees F
Module voltage high alarm threshold       : 3.630 V
Module voltage low alarm threshold        : 2.970 V
Module voltage high warning threshold     : 3.465 V
Module voltage low warning threshold      : 3.135 V
Laser rx power high alarm threshold       : 1.5849 mW / 2.00 dBm
Laser rx power low alarm threshold        : 0.0407 mW / -13.90 dBm
Laser rx power high warning threshold     : 0.7943 mW / -1.00 dBm
Laser rx power low warning threshold      : 0.1023 mW / -9.90 dBm

```

fpc4:

show virtual-chassis vc-port diagnostics optics local

```

user@switch> show virtual-chassis vc-port diagnostics optics local
Virtual chassis port: vcp-2/0
  Optical diagnostics : N/A
Virtual chassis port: vcp-2/1
  Optical diagnostics : N/A
Virtual chassis port: vcp-255/0/14
  Optical diagnostics : N/A
Virtual chassis port: vcp-255/0/15
  Optical diagnostics : N/A
Virtual chassis port: vcp-255/0/24
  Laser bias current : 4.130 mA
  Laser output power : 0.2450 mW / -6.11 dBm
  Module temperature : 32 degrees C / 90 degrees F
  Module voltage     : 3.3530 V
  Receiver signal average optical power : 0.0961 mW / -10.17 dBm
  Laser bias current high alarm : Off
  Laser bias current low alarm  : Off
  Laser bias current high warning : Off
  Laser bias current low warning : Off
  Laser output power high alarm : Off
  Laser output power low alarm  : Off
  Laser output power high warning : Off
  Laser output power low warning : Off
  Module temperature high alarm : Off
  Module temperature low alarm  : Off
  Module temperature high warning : Off
  Module temperature low warning : Off
  Module voltage high alarm : Off
  Module voltage low alarm  : Off
  Module voltage high warning : Off
  Module voltage low warning : Off
  Laser rx power high alarm : Off
  Laser rx power low alarm  : Off
  Laser rx power high warning : Off
  Laser rx power low warning : Off
  Laser bias current high alarm threshold : 14.998 mA
  Laser bias current low alarm threshold  : 0.998 mA

```



```

Laser rx power low alarm threshold      : 0.0407 mW / -13.90 dBm
Laser rx power high warning threshold   : 0.7943 mW / -1.00 dBm
Laser rx power low warning threshold    : 0.1023 mW / -9.90 dBm

```

show virtual-chassis vc-port diagnostics optics (member member-id)

```

user@switch> show virtual-chassis vc-port diagnostics optics member 2
fpc2:

```

```

-----
Virtual chassis port: vcp-2/0
  Optical diagnostics                : N/A
Virtual chassis port: vcp-2/1
  Optical diagnostics                : N/A
Virtual chassis port: vcp-255/0/14
  Optical diagnostics                : N/A
Virtual chassis port: vcp-255/0/15
  Optical diagnostics                : N/A
Virtual chassis port: vcp-255/0/24
  Laser bias current                  : 4.130 mA
  Laser output power                  : 0.2450 mW / -6.11 dBm
  Module temperature                  : 31 degrees C / 88 degrees F
  Module voltage                      : 3.3530 V
  Receiver signal average optical power : 0.0961 mW / -10.17 dBm
  Laser bias current high alarm       : Off
  Laser bias current low alarm        : Off
  Laser bias current high warning     : Off
  Laser bias current low warning      : Off
  Laser output power high alarm       : Off
  Laser output power low alarm        : Off
  Laser output power high warning     : Off
  Laser output power low warning      : Off
  Module temperature high alarm       : Off
  Module temperature low alarm        : Off
  Module temperature high warning     : Off
  Module temperature low warning      : Off
  Module voltage high alarm           : Off
  Module voltage low alarm            : Off
  Module voltage high warning         : Off
  Module voltage low warning          : Off
  Laser rx power high alarm           : Off
  Laser rx power low alarm            : Off
  Laser rx power high warning         : Off
  Laser rx power low warning          : Off
  Laser bias current high alarm threshold : 14.998 mA
  Laser bias current low alarm threshold : 0.998 mA
  Laser bias current high warning threshold : 14.000 mA
  Laser bias current low warning threshold : 1.198 mA
  Laser output power high alarm threshold : 0.7940 mW / -1.00 dBm
  Laser output power low alarm threshold : 0.0790 mW / -11.02 dBm
  Laser output power high warning threshold : 0.6300 mW / -2.01 dBm
  Laser output power low warning threshold : 0.0990 mW / -10.04 dBm
  Module temperature high alarm threshold : 85 degrees C / 185 degrees F
  Module temperature low alarm threshold : -10 degrees C / 14 degrees F
  Module temperature high warning threshold : 80 degrees C / 176 degrees F
  Module temperature low warning threshold : -5 degrees C / 23 degrees F
  Module voltage high alarm threshold : 3.600 V
  Module voltage low alarm threshold : 3.000 V
  Module voltage high warning threshold : 3.499 V
  Module voltage low warning threshold : 3.099 V
  Laser rx power high alarm threshold : 1.5848 mW / 2.00 dBm

```

```

Laser rx power low alarm threshold      : 0.0100 mW / -20.00 dBm
Laser rx power high warning threshold   : 1.2589 mW / 1.00 dBm
Laser rx power low warning threshold    : 0.0125 mW / -19.03 dBm
Virtual chassis port: vcp-255/0/3
Laser bias current                      : 5.418 mA
Laser output power                      : 0.4770 mW / -3.21 dBm
Module temperature                      : 28 degrees C / 83 degrees F
Module voltage                          : 3.3450 V
Receiver signal average optical power   : 0.3964 mW / -4.02 dBm
Laser bias current high alarm           : Off
Laser bias current low alarm            : Off
Laser bias current high warning         : Off
Laser bias current low warning          : Off
Laser output power high alarm           : Off
Laser output power low alarm            : Off
Laser output power high warning         : Off
Laser output power low warning          : Off
Module temperature high alarm           : Off
Module temperature low alarm            : Off
Module temperature high warning         : Off
Module temperature low warning          : Off
Module voltage high alarm               : Off
Module voltage low alarm                : Off
Module voltage high warning             : Off
Module voltage low warning              : Off
Laser rx power high alarm               : Off
Laser rx power low alarm                : Off
Laser rx power high warning             : Off
Laser rx power low warning              : Off
Laser bias current high alarm threshold : 10.500 mA
Laser bias current low alarm threshold  : 2.000 mA
Laser bias current high warning threshold : 9.000 mA
Laser bias current low warning threshold : 2.500 mA
Laser output power high alarm threshold : 1.4120 mW / 1.50 dBm
Laser output power low alarm threshold  : 0.0740 mW / -11.31 dBm
Laser output power high warning threshold : 0.7070 mW / -1.51 dBm
Laser output power low warning threshold : 0.1860 mW / -7.30 dBm
Module temperature high alarm threshold : 75 degrees C / 167 degrees F
Module temperature low alarm threshold  : -5 degrees C / 23 degrees F
Module temperature high warning threshold : 70 degrees C / 158 degrees F
Module temperature low warning threshold : 0 degrees C / 32 degrees F
Module voltage high alarm threshold     : 3.630 V
Module voltage low alarm threshold      : 2.970 V
Module voltage high warning threshold   : 3.465 V
Module voltage low warning threshold     : 3.135 V
Laser rx power high alarm threshold     : 1.5849 mW / 2.00 dBm
Laser rx power low alarm threshold      : 0.0407 mW / -13.90 dBm
Laser rx power high warning threshold   : 0.7943 mW / -1.00 dBm
Laser rx power low warning threshold     : 0.1023 mW / -9.90 dBm

```

test interface restart-auto-negotiation

Syntax	test interface restart-auto-negotiation <i>interface-name</i>
Release Information	Command introduced in Junos OS Release 7.6. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Restarts auto-negotiation on a Fast Ethernet or Gigabit Ethernet interface.
Options	<i>interface-name</i> —Interface name: <i>fe-fpc/pic/port</i> or <i>ge-fpc/pic/port</i> .
Required Privilege Level	view
List of Sample Output	test interface restart-auto-negotiation on page 419
Output Fields	Use the <code>show interfaces extensive</code> command to see the state for auto-negotiation.

Sample Output

test interface restart-auto-negotiation

```
user@host> test interface restart-auto-negotiation fe-1/0/0
```


PART 4

Troubleshooting

- [Troubleshooting Procedures on page 423](#)

CHAPTER 7

Troubleshooting Procedures

- [Troubleshooting an Aggregated Ethernet Interface on page 423](#)
- [Troubleshooting Network Interfaces on EX3200 Switches on page 424](#)
- [Troubleshooting Network Interfaces on EX4200 Switches on page 426](#)
- [Troubleshooting Uplink Module Installation or Replacement on EX3200 Switches on page 427](#)
- [Troubleshooting Interface Configuration and Cable Faults on page 428](#)
- [Troubleshooting Unicast RPF on page 429](#)
- [Diagnosing a Faulty Twisted-Pair Cable \(CLI Procedure\) on page 430](#)

Troubleshooting an Aggregated Ethernet Interface

Troubleshooting issues for aggregated Ethernet interfaces:

- [Show Interfaces Command Shows the LAG is Down on page 423](#)
- [Logical Interface Statistics Do Not Reflect All Traffic on page 424](#)
- [IPv6 Interface Traffic Statistics Are Not Supported on page 424](#)
- [SNMP Counters ifHCInBroadcastPkts and ifInBroadcastPkts Are Always 0 on page 424](#)

Show Interfaces Command Shows the LAG is Down

Problem **Description:** The `show interfaces terse` command shows that the LAG is down.

Solution Check the following:

- Verify that there is no configuration mismatch.
- Verify that all member ports are up.
- Verify that a LAG is part of family ethernet—switching (Layer 2 LAG) or family inet (Layer 3 LAG).
- Verify that the LAG member is connected to the correct LAG at the other end.
- Verify that the LAG members belong to the same switch (or the same Virtual Chassis).

Logical Interface Statistics Do Not Reflect All Traffic

Problem **Description:** The traffic statistics for a logical interface do not include all of the traffic.

Solution Traffic statistics fields for logical interfaces in **show interfaces** commands show only control traffic; the traffic statistics do not include data traffic. You can view the statistics for all traffic only per physical interface.

IPv6 Interface Traffic Statistics Are Not Supported

Problem **Description:** The IPv6 transit statistics in the **show interfaces** command display all 0 values.

Solution EX Series switches do not support the collection and reporting of IPv6 transit statistics.

SNMP Counters ifHCInBroadcastPkts and ifInBroadcastPkts Are Always 0

Problem **Description:** The values for the SNMP counters ifHCInBroadcastPkts and ifInBroadcastPkts are always 0.

Solution The SNMP counters ifHCInBroadcastPkts and ifInBroadcastPkts are not supported for aggregated Ethernet interfaces on EX Series switches.

Related Documentation

- [Verifying the Status of a LAG Interface on page 279](#)
- [Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch on page 31](#)
- [Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch on page 37](#)

Troubleshooting Network Interfaces on EX3200 Switches

This topic provides troubleshooting information for specific problems related to interfaces on EX3200 switches.

- [The interface on one of the last four built-in network ports in an EX3200 switch \(for example, interface ge-0/0/23\) is down on page 425](#)
- [The interface on the port in which an SFP or SFP+ transceiver is installed in an SFP+ uplink module is down on page 425](#)

The interface on one of the last four built-in network ports in an EX3200 switch (for example, interface ge-0/0/23) is down

Problem Description: The interface on one of the last four built-in ports (ge-0/0/20 through ge-0/0/23 on 24-port models or ge-0/0/44 through ge-0/0/47 on 48-port models) of an EX3200 switch is down.

Environment: An SFP or SFP+ uplink module is installed in the switch and a transceiver is installed in one of the ports on the uplink module.

Symptoms: When you check the status with the CLI command `show interfaces ge-` or with the J-Web user interface, the disabled port is not listed.

Cause The last four built-in ports use the same ASIC as the SFP uplink module. Therefore, if you install a transceiver in an SFP or SFP+ uplink module installed in an EX3200 switch, a corresponding base port from the last four built-in ports is disabled.

Solution If you need to use the disabled built-in port, you must remove the transceiver from the SFP or SFP+ uplink module. Alternatively, you can install an XFP uplink module instead of an SFP or SFP+ uplink module. There is no conflict between the built-in network ports and the ports on the XFP uplink modules.

The interface on the port in which an SFP or SFP+ transceiver is installed in an SFP+ uplink module is down

Problem Description: The interface on the port in which an SFP or SFP+ transceiver is installed in an SFP+ uplink module installed in an EX3200 switch is down.

Symptoms: When you check the status with the CLI command `show interfaces ge-` or with the J-Web user interface, the disabled port is not listed.

Cause By default, the SFP+ uplink module operates in the 10-gigabit mode and supports only SFP+ transceivers. The operating mode for the module is incorrectly set.

Solution Either SFP+ or SFP transceivers can be installed in SFP+ uplink modules. You must configure the operating mode of the SFP+ uplink module to match the type of transceiver you want to use. For SFP+ transceivers, configure the 10-gigabit operating mode and for SFP transceivers, configure the 1-gigabit operating mode. See [“Setting the Mode on an SFP+ or SFP+ MACSec Uplink Module \(CLI Procedure\)” on page 129](#).

Related Documentation

- [Troubleshooting Uplink Module Installation or Replacement on EX3200 Switches on page 427](#)
- [Monitoring Interface Status and Traffic on page 277](#)
- [Configuring Gigabit Ethernet Interfaces \(CLI Procedure\) on page 64](#)

- [Configuring Gigabit Ethernet Interfaces \(J-Web Procedure\) on page 68](#)
- [Removing a Transceiver from a Switch](#)
- [Uplink Modules in EX3200 Switches](#)
- [EX Series Switches Interfaces Overview on page 3](#)

Troubleshooting Network Interfaces on EX4200 Switches

This topic provides troubleshooting information for specific problems related to interfaces on EX4200 switches.

- [The interface on the port in which an SFP or SFP+ transceiver is installed is down on page 426](#)

The interface on the port in which an SFP or SFP+ transceiver is installed is down

Problem **Description:** The interface on the port in which an SFP or SFP+ transceiver is installed in an uplink module installed in an EX4200 switch is down.

Symptoms: When you check the status with the CLI command [show interfaces ge-](#) or with the J-Web user interface, the disabled port is not listed.

Cause By default, the SFP+ and SFP+ MACsec uplink modules operate in the 10-gigabit mode and support only SFP+ transceivers. The operating mode for the module is incorrectly set.

Solution Either SFP+ or SFP transceivers can be installed in the uplink modules. You must configure the operating mode of the SFP+ or SFP+ MACsec uplink module to match the type of transceiver you want to use. For SFP+ transceivers, configure the 10-gigabit operating mode and for SFP transceivers, configure the 1-gigabit operating mode. See [“Setting the Mode on an SFP+ or SFP+ MACSec Uplink Module \(CLI Procedure\)” on page 129](#).

- Related Documentation**
- [Troubleshooting Virtual Chassis Port Connectivity on an EX4200 Switch](#)
 - [Monitoring Interface Status and Traffic on page 277](#)
 - [Configuring Gigabit Ethernet Interfaces \(CLI Procedure\) on page 64](#)
 - [Configuring Gigabit Ethernet Interfaces \(J-Web Procedure\) on page 68](#)
 - [Removing a Transceiver from a Switch](#)
 - [Uplink Modules in EX4200 Switches](#)
 - [EX Series Switches Interfaces Overview on page 3](#)

Troubleshooting Uplink Module Installation or Replacement on EX3200 Switches

This topic provides troubleshooting information for specific problems related to uplink module ports on EX3200 switches.

1. [One of the last four network ports on an EX3200 switch with an SFP or SFP+ uplink module installed is disabled on page 427](#)
2. [A port on an SFP uplink module installed in an EX3200 switch is disabled on page 427](#)

One of the last four network ports on an EX3200 switch with an SFP or SFP+ uplink module installed is disabled

Problem Description: One of the last four built-in ports (**ge-0/0/20** through **ge-0/0/23** on 24-port models or **ge-0/0/44** through **ge-0/0/47** on 48-port models) of an EX3200 switch with an SFP or SFP+ uplink module installed in it is disabled.

Symptoms: When you check the status with the CLI command **show interfaces ge-** or with the J-Web user interface, the disabled port is not listed.

Cause The last four built-in ports use the same ASIC as the SFP uplink module. Therefore, if you install a transceiver in an SFP or SFP+ uplink module installed in an EX3200 switch, a corresponding base port from the last four built-in ports is disabled.

Solution If you need to use the disabled built-in port, you must remove the transceiver from the SFP or SFP+ uplink module. Alternatively, you can install an XFP uplink module instead of an SFP or SFP+ uplink module. There is no conflict between the built-in network ports and the ports on the XFP uplink modules.

A port on an SFP uplink module installed in an EX3200 switch is disabled

Problem Description: One of the ports (**ge-0/1/0** through **ge-0/1/3**) of an SFP uplink module installed in an EX3200 switch is disabled.

Symptoms: When you check the status with the **show interface** commands or with the J-Web user interface, the disabled port is not listed.

Cause If you replace a transceiver multiple times in quick succession in a port in an SFP uplink module installed in an EX3200 switch, it might cause an eeprom read problem. The switch might not create an interface for that port and that port might be disabled.

Solution To enable the disabled uplink module port, remove the transceiver from that port and install it after 10 seconds.

Related Documentation

- [Monitoring Interface Status and Traffic on page 277](#)
- [Configuring Gigabit Ethernet Interfaces \(CLI Procedure\) on page 64](#)

