




Junos[®] OS for EX Series Ethernet Switches, Release 11.4: Interfaces



Published: 2011-11-16
Revision 1

Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, California 94089
USA
408-745-2000
www.juniper.net

This product includes the Envoy SNMP Engine, developed by Epilogue Technology, an Integrated Systems Company. Copyright © 1986-1997, Epilogue Technology Corporation. All rights reserved. This program and its documentation were developed at private expense, and no part of them is in the public domain.

This product includes memory allocation software developed by Mark Moraes, copyright © 1988, 1989, 1993, University of Toronto.

This product includes FreeBSD software developed by the University of California, Berkeley, and its contributors. All of the documentation and software included in the 4.4BSD and 4.4BSD-Lite Releases is copyrighted by the Regents of the University of California. Copyright © 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994. The Regents of the University of California. All rights reserved.

GateD software copyright © 1995, the Regents of the University. All rights reserved. Gate Daemon was originated and developed through release 3.0 by Cornell University and its collaborators. Gated is based on Kirton's EGP, UC Berkeley's routing daemon (routed), and DCN's HELLO routing protocol. Development of Gated has been supported in part by the National Science Foundation. Portions of the GateD software copyright © 1988, Regents of the University of California. All rights reserved. Portions of the GateD software copyright © 1991, D. L. S. Associates.

This product includes software developed by Maker Communications, Inc., copyright © 1996, 1997, Maker Communications, Inc.

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.

Products made or sold by Juniper Networks or components thereof might be covered by one or more of the following patents that are owned by or licensed to Juniper Networks: U.S. Patent Nos. 5,473,599, 5,905,725, 5,909,440, 6,192,051, 6,333,650, 6,359,479, 6,406,312, 6,429,706, 6,459,579, 6,493,347, 6,538,518, 6,538,899, 6,552,918, 6,567,902, 6,578,186, and 6,590,785.

Junos® OS for EX Series Ethernet Switches, Release 11.4: Interfaces

Copyright © 2011, Juniper Networks, Inc.

All rights reserved.

Revision History

November 2011—Revision 1

The information in this document is current as of the date listed in the revision history.

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 This Topic Collection	ix
	How to Use This Guide	ix
	List of EX Series Guides for Junos OS Release 11.4	ix
	Downloading Software	xi
	Documentation Symbols Key	xii
	Documentation Feedback	xiii
	Requesting Technical Support	xiv
	Self-Help Online Tools and Resources	xiv
	Opening a Case with JTAC	xiv
Part 1	Interfaces on EX Series Switches	
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	7
	Wildcard Characters in Interface Names	7
	Understanding Aggregated Ethernet Interfaces and LACP	8
	Link Aggregation Group (LAG)	8
	Link Aggregation Control Protocol (LACP)	9
	Understanding Interface Ranges on EX Series Switches	10
	Understanding Layer 3 Subinterfaces	12
	Understanding Unicast RPF for EX Series Switches	13
	Unicast RPF for EX Series Switches Overview	13
	Unicast RPF Implementation for EX Series Switches	14
	Unicast RPF Packet Filtering	14
	Bootstrap Protocol (BOOTP) and DHCP Requests	14
	Default Route Handling	14
	When to Enable Unicast RPF	14
	When Not to Enable Unicast RPF	15
	Limitations of the Unicast RPF Implementation on EX3200 and EX4200 Switches	16
	Understanding IP Directed Broadcast for EX Series Switches	17
	IP Directed Broadcast for EX Series Switches Overview	17
	IP Directed Broadcast Implementation for EX Series Switches	18
	When to Enable IP Directed Broadcast	18
	When Not to Enable IP Directed Broadcast	18
	802.1Q VLANs Overview	19

Chapter 2	Examples of Interfaces Configuration	21
	Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch	21
	Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch	27
	Example: Configuring Layer 3 Subinterfaces for a Distribution Switch and an Access Switch	32
	Example: Configuring Unicast RPF on an EX Series Switch	39
	Example: Configuring IP Directed Broadcast on an EX Series Switch	43
Chapter 3	Configuring Interfaces	47
	Configuring Gigabit Ethernet Interfaces (CLI Procedure)	48
	Configuring VLAN Options and Port Mode	48
	Configuring the Link Settings	49
	Configuring the IP Options	51
	Configuring Gigabit Ethernet Interfaces (J-Web Procedure)	51
	Port Role Configuration with the J-Web Interface (with CLI References)	58
	Adding an Interface Description to the Configuration	62
	Example: Adding an Interface Description to the Configuration	62
	Adding a Logical Unit Description to the Configuration	63
	Disabling a Physical Interface	64
	Example: Disabling a Physical Interface	64
	Disabling a Logical Interface	65
	Configuring Flow Control	65
	Configuring the Interface Address	66
	Configuring Interface IPv4 Addresses	67
	Configuring Interface IPv6 Addresses	68
	Configuring the Interface Bandwidth	68
	Configuring the Media MTU	69
	Setting the Protocol MTU	80
	Interface Ranges	80
	Configuring Interface Ranges	81
	Expanding Interface Range Member and Member Range Statements	84
	Configuration Inheritance for Member Interfaces	86
	Member Interfaces Inheriting Configuration from Configuration Groups	87
	Interfaces Inheriting Common Configuration	88
	Configuring Inheritance Range Priorities	88
	Configuration Expansion Where Interface Range Is Used	89
	Configuring Accounting for the Physical Interface	90
	Applying an Accounting Profile to the Physical Interface	90
	Example: Applying an Accounting Profile to the Physical Interface	90
	Configuring Accounting for the Logical Interface	91
	Applying an Accounting Profile to the Logical Interface	91
	Example: Applying an Accounting Profile to the Logical Interface	91
	Configuring Ethernet Loopback Capability	92
	Configuring Gratuitous ARP	93

	Configuring Static ARP Table Entries	94
	Example: Configuring Static ARP Table Entries	95
	Disabling the Transmission of Redirect Messages on an Interface	95
	Configuring Restricted and Unrestricted Proxy ARP	95
	Enabling or Disabling SNMP Notifications on Logical Interfaces	96
	Enabling or Disabling SNMP Notifications on Physical Interfaces	97
	Configuring Aggregated Ethernet Interfaces (CLI Procedure)	97
	Configuring Aggregated Ethernet Interfaces (J-Web Procedure)	98
	Configuring Aggregated Ethernet LACP (CLI Procedure)	101
	Configuring Aggregated Ethernet Link Protection	102
	Configuring Link Protection for Aggregated Ethernet Interfaces	102
	Configuring Primary and Backup Links for Link Aggregated Ethernet Interfaces	102
	Reverting Traffic to a Primary Link When Traffic is Passing Through a Backup Link	103
	Disabling Link Protection for Aggregated Ethernet Interfaces	103
	Configuring Aggregated Ethernet Link Speed	103
	Configuring Aggregated Ethernet Minimum Links	104
	Configuring Tagged Aggregated Ethernet Interfaces	105
	Configuring a Layer 3 Subinterface (CLI Procedure)	105
	Configuring Unicast RPF (CLI Procedure)	106
	Disabling Unicast RPF (CLI Procedure)	107
	Configuring IP Directed Broadcast (CLI Procedure)	108
	Tracing Operations of an Individual Router or Switch Interface	109
	Tracing Operations of the Interface Process	110
	Setting the Mode on an SFP+ Uplink Module (CLI Procedure)	110
	Configuring the Media Type on Dual-Purpose Uplink Ports (CLI Procedure)	111
Chapter 4	Verifying Interfaces	113
	Monitoring Interface Status and Traffic	113
	Verifying the Status of a LAG Interface	114
	Verifying That LACP Is Configured Correctly and Bundle Members Are Exchanging LACP Protocol Packets	115
	Verifying the LACP Setup	115
	Verifying That LACP Packets Are Being Exchanged	115
	Verifying That Layer 3 Subinterfaces Are Working	116
	Verifying Unicast RPF Status	117
	Verifying IP Directed Broadcast Status	119
Chapter 5	Troubleshooting Interfaces	121
	Troubleshooting Network Interfaces on EX3200 Switches	121
	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	121
	The interface on the port in which an SFP or SFP+ transceiver is installed in an SFP+ uplink module is down	122
	Troubleshooting Network Interfaces on EX4200 Switches	122
	The interface on the port in which an SFP or SFP+ transceiver is installed in an SFP+ uplink module is down	122

	Troubleshooting an Aggregated Ethernet Interface	123
	Show Interfaces Command Shows the LAG is Down	123
	Logical Interface Statistics Do Not Reflect All Traffic	123
	Troubleshooting Interface Configuration and Cable Faults	124
	Interface Configuration or Connectivity Is Not Working	124
	Troubleshooting Unicast RPF	125
	Legitimate Packets Are Discarded	125
	Troubleshooting Virtual Chassis Port Connectivity on an EX4200 Switch	125
	Virtual Chassis port (VCP) connection does not work	126
	Diagnosing a Faulty Twisted-Pair Cable (CLI Procedure)	126
Chapter 6	Configuration Statements for Interfaces	129
	[edit chassis] Configuration Statement Hierarchy	129
	[edit interfaces] Configuration Statement Hierarchy	130
	802.3ad	135
	accounting-profile	136
	address	137
	aggregated-devices	139
	aggregated-ether-options	140
	arp	141
	auto-negotiation	142
	bandwidth	144
	broadcast	145
	chassis	146
	description	147
	device-count	148
	disable (Interface)	149
	ether-options	150
	ethernet	151
	eui-64	151
	family (for EX Series switches)	152
	filter	156
	flow-control	157
	force-up	157
	gratuitous-arp-reply	158
	interface-range	159
	interfaces (for EX Series switches)	161
	lACP (802.3ad)	168
	lACP (Aggregated Ethernet)	169
	link-mode	170
	link-protection	171
	link-speed (Aggregated Ethernet)	172
	loopback (Aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet)	173
	media-type	174
	member	174
	members	175
	member-range	177
	minimum-links	178
	mtu	179

	native-vlan-id	180
	no-redirects	181
	periodic	182
	pic	183
	pic-mode	183
	port-mode	184
	preferred	185
	primary (Address on Interface)	186
	proxy-arp	187
	rpf-check	188
	sfpplus	189
	speed	190
	targeted-broadcast	191
	traceoptions (Individual Interfaces)	192
	traceoptions (Interface Process)	194
	traps	195
	unit	196
	vlan	197
	vlan-id	198
	vlan-tagging	199
Chapter 7	Operational Commands for Interfaces	201
	clear ipv6 neighbors	202
	monitor interface	203
	request diagnostics tdr	210
	show diagnostics tdr	212
	show ethernet-switching interfaces	217
	show interfaces diagnostics optics	221
	show interfaces ge	228
	show interfaces me0	239
	show interfaces queue	246
	show interfaces vlan	252
	show interfaces xe-	264
	show ipv6 neighbors	277
	show lacp interfaces	279
	test interface restart-auto-negotiation	284

About This Topic Collection

- [How to Use This Guide on page ix](#)
- [List of EX Series Guides for Junos OS Release 11.4 on page ix](#)
- [Downloading Software on page xi](#)
- [Documentation Symbols Key on page xii](#)
- [Documentation Feedback on page xiii](#)
- [Requesting Technical Support on page xiv](#)

How to Use This Guide

Complete documentation for the EX Series product family is provided on webpages at http://www.juniper.net/techpubs/en_US/release-independent/information-products/pathway-pages/ex-series/product/index.html. We have selected content from these webpages and created a number of EX Series guides that collect related topics into a book-like format so that the information is easy to print and easy to download to your local computer.

Software features for EX Series switches are listed by platform and by Junos OS release in a standalone document. See [EX Series Switch Software Features Overview](#).

The release notes are at http://www.juniper.net/techpubs/en_US/junos11.4/information-products/topic-collections/release-notes/11.4/junos-release-notes-11.4.pdf.

List of EX Series Guides for Junos OS Release 11.4

Title	Description
<i>Complete Hardware Guide for EX2200 Ethernet Switches</i>	Component descriptions, site preparation, installation, replacement, and safety and compliance information for EX2200 Ethernet switches
<i>Complete Hardware Guide for EX3200 Ethernet Switches</i>	Component descriptions, site preparation, installation, replacement, and safety and compliance information for EX3200 Ethernet switches
<i>Complete Hardware Guide for EX3300 Ethernet Switches</i>	Component descriptions, site preparation, installation, replacement, and safety and compliance information for EX3300 Ethernet switches

Title	Description
<i>Complete Hardware Guide for EX4200 Ethernet Switches</i>	Component descriptions, site preparation, installation, replacement, and safety and compliance information for EX4200 Ethernet switches
<i>Complete Hardware Guide for EX4500 Ethernet Switches</i>	Component descriptions, site preparation, installation, replacement, and safety and compliance information for EX4500 Ethernet switches
<i>Complete Hardware Guide for EX6210 Ethernet Switches</i>	Component descriptions, site preparation, installation, replacement, and safety and compliance information for EX6210 Ethernet switches
<i>Complete Hardware Guide for EX8208 Ethernet Switches</i>	Component descriptions, site preparation, installation, replacement, and safety and compliance information for EX8208 Ethernet switches
<i>Complete Hardware Guide for EX8216 Ethernet Switches</i>	Component descriptions, site preparation, installation, replacement, and safety and compliance information for EX8216 Ethernet switches
<i>Complete Hardware Guide for the XRE200 External Routing Engine</i>	Component descriptions, site preparation, installation, replacement, and safety and compliance information for the XRE200 External Routing Engine
<i>Complete Software Guide for Junos® OS for EX Series Ethernet Switches, Release 11.4</i>	Software feature descriptions, configuration examples, and tasks for Junos OS for EX Series switches
Software Topic Collections	Software feature descriptions, configuration examples and tasks, and reference pages for configuration statements and operational commands (This information also appears in the <i>Complete Software Guide for Junos® OS for EX Series Ethernet Switches, Release 11.4.</i>)
<i>Junos® OS for EX Series Ethernet Switches, Release 11.4: Access and User Management</i>	
<i>Junos® OS for EX Series Ethernet Switches, Release 11.4: Access Control</i>	
<i>Junos® OS for EX Series Ethernet Switches, Release 11.4: Configuration Management</i>	
<i>Junos® OS for EX Series Ethernet Switches, Release 11.4: Class of Service</i>	
<i>Junos® OS for EX Series Ethernet Switches, Release 11.4: Device Security</i>	
<i>Junos® OS for EX Series Ethernet Switches, Release 11.4: Ethernet Switching</i>	
<i>Junos® OS for EX Series Ethernet Switches, Release 11.4: EX3300, EX4200, and EX4500 Virtual Chassis</i>	





Title	Description
<i>Junos® OS for EX Series Ethernet Switches, Release 11.4: EX8200 Virtual Chassis</i>	
<i>Junos® OS for EX Series Ethernet Switches, Release 11.4: Fibre Channel over Ethernet</i>	
<i>Junos® OS for EX Series Ethernet Switches, Release 11.4: High Availability</i>	
<i>Junos® OS for EX Series Ethernet Switches, Release 11.4: Interfaces</i>	
<i>Junos® OS for EX Series Ethernet Switches, Release 11.4: Layer 3 Protocols</i>	
<i>Junos® OS for EX Series Ethernet Switches, Release 11.4: MPLS</i>	
<i>Junos® OS for EX Series Ethernet Switches, Release 11.4: Multicast</i>	
<i>Junos® OS for EX Series Switches, Release 11.4: Network Management and Monitoring</i>	
<i>Junos® OS for EX Series Switches, Release 11.4: Port Security</i>	
<i>Junos® OS for EX Series Switches, Release 11.4: Power over Ethernet</i>	
<i>Junos® OS for EX Series Ethernet Switches, Release 11.4: Routing Policy and Packet Filtering</i>	
<i>Junos® OS for EX Series Ethernet Switches, Release 11.4: Software Installation</i>	
<i>Junos® OS for EX Series Ethernet Switches, Release 11.4: Spanning-Tree Protocols</i>	
<i>Junos® OS for EX Series Ethernet Switches, Release 11.4: System Monitoring</i>	
<i>Junos® OS for EX Series Ethernet Switches, Release 11.4: System Services</i>	
<i>Junos® OS for EX Series Ethernet Switches, Release 11.4: System Setup</i>	
<i>Junos® OS for EX Series Ethernet Switches, Release 11.4: User Interfaces</i>	

Downloading Software

You can download Junos OS for EX Series switches from the Download Software area at <http://www.juniper.net/customers/support/> . To download the software, you must

have a Juniper Networks user account. For information about obtaining an account, see <http://www.juniper.net/entitlement/setupAccountInfo.do>.

Documentation Symbols Key

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.
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
Fixed-width text like this	Represents output that appears on the terminal screen.	user@host> show chassis alarms No alarms currently active
<i>Italic text like this</i>	<ul style="list-style-type: none"> Introduces important new terms. Identifies book 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 System Basics Configuration 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>
Plain text like this	Represents names of configuration statements, commands, files, and directories; IP addresses; 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.

Text and Syntax Conventions		
Convention	Description	Examples
< > (angle brackets)	Enclose 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)	Enclose a variable for which you can substitute one or more values.	community name members [<i>community-ids</i>]
Indentation and braces ({ })	Identify 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.	
J-Web GUI Conventions		
Bold text like this	Represents J-Web 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.
> (bold right angle bracket)	Separates levels in a hierarchy of J-Web 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. Send e-mail to techpubs-comments@juniper.net with the following:

- Document URL or title
- Page number if applicable
- Software version
- Your name and company

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 JNASC 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: <https://www.juniper.net/alerts/>
- 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

Interfaces on EX Series Switches

- [Interfaces—Overview on page 3](#)
- [Examples of Interfaces Configuration on page 21](#)
- [Configuring Interfaces on page 47](#)
- [Verifying Interfaces on page 113](#)
- [Troubleshooting Interfaces on page 121](#)
- [Configuration Statements for Interfaces on page 129](#)
- [Operational Commands for Interfaces on page 201](#)

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 Interface Ranges on EX Series Switches on page 10](#)
- [Understanding Layer 3 Subinterfaces on page 12](#)
- [Understanding Unicast RPF for EX Series Switches on page 13](#)
- [Understanding IP Directed Broadcast for EX Series Switches on page 17](#)
- [802.1Q VLANs Overview on page 19](#)

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 on these interfaces. For additional information, see the [Junos OS Interfaces Fundamentals Configuration Guide](#).

For information on 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 1 on page 3](#) lists the types of network interfaces supported on EX Series switches.

Table 1: 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.

Table 1: Network Interface Types and Purposes (*continued*)

Type	Purpose
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 port 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 2 on page 4 lists the types of special interfaces supported on EX Series switches.

Table 2: 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 on the console port in a Virtual Chassis, see Understanding Global Management of an EX3300, EX4200, or EX4500 Virtual Chassis.
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.)
Routed VLAN Interface (RVI)	EX Series switches use a Layer 3 routed VLAN interface (RVI) named vlan 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. The RVI functions as a logical router, eliminating the need for having both a switch and a router. The RVI (the vlan interface) must be configured as part of a broadcast domain or virtual private LAN service (VPLS) routing instance for Layer 3 traffic to be routed out of it.

Table 2: Special Interface Types and Purposes (*continued*)

Type	Purpose
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 Setting an Uplink Port on an EX3300 or EX4200 Switch as a Virtual Chassis Port (CLI Procedure). 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 Understanding the High-Speed Interconnection of the EX4200 and EX4500 Virtual Chassis Members. <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 Setting an Uplink Port on an EX3300 or EX4200 Switch as a Virtual Chassis Port (CLI Procedure) or Setting an SFP+ Port as a Virtual Chassis Port on an EX4500 Switch (CLI Procedure).</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 Understanding Virtual Chassis Ports in an EX8200 Virtual Chassis.</p>
Virtual management Ethernet (VME) interface	<p>EX3300, EX4200, 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 on the VME interface, see Understanding Global Management of an EX3300, EX4200, or EX4500 Virtual Chassis.</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](#)
- [EX4500 Switches Hardware Overview](#)
- [EX6210 Switch Hardware Overview](#)
- [EX8208 Switch Hardware Overview](#)
- [EX8216 Switch Hardware Overview](#)
- [XRE200 External Routing Engine Hardware Overview](#)

- [PoE and EX Series Switches Overview](#)
- [Understanding Aggregated Ethernet Interfaces and LACP on page 8](#)
- [Understanding Layer 3 Subinterfaces on page 12](#)

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 on 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 7](#)
- [Wildcard Characters in Interface Names on page 7](#)

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
- *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, 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 in the Virtual Chassis.
 - On an EX6200 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, and EX4500 switches, the PIC number is **0** for all built-in interfaces (interfaces that are not an uplink port).
 - On EX2200, EX3200, EX3300, 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 right-hand uplink module.
 - 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, and EX4500 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, and EX4500 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](#)
 - Front Panel of an EX2200 Switch
 - Front Panel of an EX3200 Switch
 - Front Panel of an EX3300 Switch
 - Front Panel of an EX4200 Switch
 - Front Panel of an EX4500 Switch
 - Slot Numbering for an EX6210 Switch
 - Slot Numbering for an EX8208 Switch
 - Slot Numbering for an EX8216 Switch

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 8](#)
- [Link Aggregation Control Protocol \(LACP\) on page 9](#)

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 3 on page 9](#) lists the EX Series switches and the maximum number of interfaces per LAG and maximum number of LAGs they support.

Table 3: 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	32
EX4200 and EX4200 Virtual Chassis	8	64
EX4500 and EX4500 Virtual Chassis	8	64
EX6200	8	64
EX8200	12	255
EX8200 Virtual Chassis	12	239

When configuring LAGs, consider the following guidelines:

- The LAG must be configured on both sides of the link.
- The interfaces on either side of the link must be set to the same speed.
- You can configure and apply firewall filters on a LAG.
- LACP can optionally be configured for link negotiation.

You can combine physical Ethernet ports belonging to different member switches of a Virtual Chassis configuration to form a LAG. See *Understanding EX3300, EX4200, and EX4500 Virtual Chassis Link Aggregation* and *Understanding Link Aggregation into an EX8200 Virtual Chassis*.



NOTE: The interfaces that are included within a bundle or 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 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.

About enabling LACP:

- 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 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. 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 may 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 there are no received PDUs, 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 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.

Related Documentation

- [Understanding EX3300, EX4200, and EX4500 Virtual Chassis Link Aggregation](#)
- [Understanding Link Aggregation into an EX8200 Virtual Chassis](#)
- [Understanding Redundant Trunk Links on EX Series Switches](#)
- [Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch on page 21](#)
- [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 27](#)
- [Junos OS Network Interfaces Configuration Guide](#)

Understanding Interface Ranges on EX Series Switches

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 bpd-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 80](#)
- [EX Series Switches Interfaces Overview on page 3](#)
- [Configuring Gigabit Ethernet Interfaces \(CLI Procedure\) on page 48](#)
- [Configuring Aggregated Ethernet Interfaces \(CLI Procedure\) on page 97](#)
- [Configuring a Layer 3 Subinterface \(CLI Procedure\) on page 105](#)
- [interface-range on page 159](#)

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 32](#)
 - [Junos OS Ethernet Interfaces Configuration Guide](#)

Understanding Unicast RPF for EX Series Switches

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, 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.



NOTE: On Juniper Networks EX3200 and EX4200 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 and EX4200 Switches” on page 16](#).

This topic covers:

- [Unicast RPF for EX Series Switches Overview on page 13](#)
- [Unicast RPF Implementation for EX Series Switches on page 14](#)
- [When to Enable Unicast RPF on page 14](#)
- [When Not to Enable Unicast RPF on page 15](#)
- [Limitations of the Unicast RPF Implementation on EX3200 and EX4200 Switches on page 16](#)

Unicast RPF for EX Series 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. Strict mode unicast RPF recognizes only one best return path to a unicast source address.

Use strict mode unicast RPF only on symmetrically routed interfaces. (For information about symmetrically routed interfaces, see [“When to Enable Unicast RPF” on page 14.](#))

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 for EX Series Switches

This section includes:

- [Unicast RPF Packet Filtering on page 14](#)
- [Bootstrap Protocol \(BOOTP\) and DHCP Requests on page 14](#)
- [Default Route Handling on page 14](#)

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

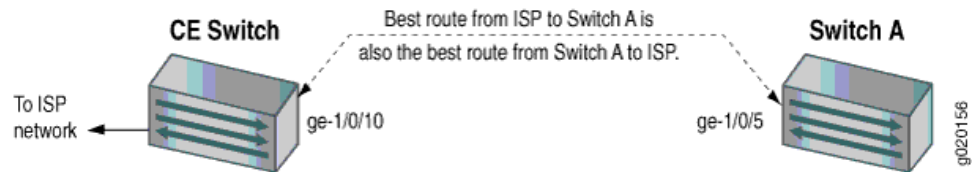
If the best return path to the source is the default route (0.0.0.0) and the default route points to **reject**, the switch discards all unicast RPF packets. If the default route points to a valid network interface, the switch performs a normal unicast RPF check on the packets.

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 15](#). 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 and EX4200 switches, ensure that *all* interfaces are symmetrically routed before you enable unicast RPF on those 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 16](#). 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 and EX4200 switches if any switch interfaces are asymmetrically routed, because unicast RPF is enabled globally on all interfaces of those 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 and EX4200 Switches

On EX3200 and EX4200 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) 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.

The 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 EX3200 and EX4200 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 39](#)
- [Configuring Unicast RPF \(CLI Procedure\) on page 106](#)
- [Disabling Unicast RPF \(CLI Procedure\) on page 107](#)

Understanding IP Directed Broadcast for EX Series Switches

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 for EX Series Switches Overview on page 17](#)
- [IP Directed Broadcast Implementation for EX Series Switches on page 18](#)
- [When to Enable IP Directed Broadcast on page 18](#)
- [When Not to Enable IP Directed Broadcast on page 18](#)

IP Directed Broadcast for EX Series Switches 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 for EX Series Switches

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 an EX Series Switch on page 43](#)
- [Configuring IP Directed Broadcast \(CLI Procedure\) on page 108](#)

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 Encapsulation](#)
- [Configuring 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 a Logical Interface for Access Mode](#)
- [Configuring a Logical Interface for Trunk Mode](#)
- [Configuring the VLAN ID List for a Trunk Interface](#)
- [Configuring a Trunk Interface on a Bridge Network](#)
- [Junos OS Ethernet Interfaces Configuration Guide](#)

CHAPTER 2

Examples of Interfaces Configuration

- [Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch on page 21](#)
- [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 27](#)
- [Example: Configuring Layer 3 Subinterfaces for a Distribution Switch and an Access Switch on page 32](#)
- [Example: Configuring Unicast RPF on an EX Series Switch on page 39](#)
- [Example: Configuring IP Directed Broadcast on an EX Series Switch on page 43](#)

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 21](#)
- [Overview and Topology on page 22](#)
- [Configuration on page 24](#)
- [Verification on page 26](#)
- [Troubleshooting on page 27](#)

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 or EX4500 Virtual Chassis \(CLI Procedure\)](#).
- Configured the uplink ports on the switches as trunk ports. See [“Configuring Gigabit Ethernet Interfaces \(CLI Procedure\)” on page 48](#).

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




Table 4 on page 23 details the topology used in this configuration example.

Table 4: 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 VCID 1	EX4200-48P switch	One XFP uplink module	0	xe-0/1/0 to SWD-0 xe-0/1/1 to SWD-1
SWA-1	Host-A Access switch VCID 1	EX4200-48P switch	One XFP uplink module	1	xe-1/1/0 to SWD-0 xe-1/1/1 to SWD-1
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

Table 4: 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-1	Host-D Distribution switch	EX4200 L-24F switch	One XFP uplink module	1	xe-1/1/0 to SWA-0
	VCID 4				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
```
- 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 26](#)
- [Verifying That LAG ae1 Has Been Created on page 26](#)

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"> • Example: Configuring an EX4200 Virtual Chassis with a Master and Backup in a Single Wiring Closet • 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 27 • 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 • Uplink Modules in EX4200 Switches

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 21:

- [Requirements on page 28](#)
- [Overview and Topology on page 28](#)
- [Configuring LACP for the LAGs on the Virtual Chassis Access Switch on page 28](#)
- [Configuring LACP for the LAGs on the Virtual Chassis Distribution Switch on page 29](#)

- [Verification on page 30](#)
- [Troubleshooting on page 31](#)

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 or EX4500 Virtual Chassis \(CLI Procedure\)](#).
- Configured the uplink ports on the switches as trunk ports. See [“Configuring Gigabit Ethernet Interfaces \(CLI Procedure\)” on page 48](#).
- 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 21](#).

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 21](#). 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 30](#)
- [Verifying That the LACP Packets Are Being Exchanged on page 30](#)

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

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 32](#)
- [Overview and Topology on page 32](#)
- [Configuring the Access Switch Subinterfaces on page 33](#)
- [Configuring the Distribution Switch Subinterfaces on page 35](#)
- [Verification on page 37](#)

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 5 on page 33](#) lists the settings for the example topology.

Table 5: 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 1.1.1/24
set interfaces ge-0/1/0 unit 1 vlan-id 102 family inet address 2.1.1/24
set interfaces ge-0/1/0 unit 2 vlan-id 103 family inet address 3.1.1/24
set interfaces ge-0/1/0 unit 3 vlan-id 104 family inet address 4.1.1/24
set interfaces ge-0/1/0 unit 4 vlan-id 105 family inet address 5.1.1/24
```

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;  
      }  
    }  
  }  
}
```



```

    }
  }
  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;
  }
}

```

Verification

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

- [Verifying That Subinterfaces Were Created on page 37](#)
- [Verifying That Traffic Passes Between VLANs on page 37](#)

Verifying That Subinterfaces Were Created

Purpose Verify that the subinterfaces were properly created on the access switch and distribution switch.

Action 1. Use the **show interfaces** command on the access switch:

```
user@access-switch> show interfaces ge-0/1/0 terse
```

Interface	Admin	Link	Proto	Local	Remote
ge-0/1/0	up	up			
ge-0/1/0.0	up	up	inet	1.1.1.1/24	
ge-0/1/0.1	up	up	inet	2.1.1.1/24	
ge-0/1/0.2	up	up	inet	3.1.1.1/24	
ge-0/1/0.3	up	up	inet	4.1.1.1/24	
ge-0/1/0.4	up	up	inet	5.1.1.1/24	
ge-0/1/0.32767	up	up			

2. Use the **show interfaces** command on the distribution switch:

```
user@distribution-switch> show interfaces ge-0/0/0 terse
```

Interface	Admin	Link	Proto	Local	Remote
ge-0/0/0	up	up			
ge-0/0/0.0	up	up	inet	1.1.1.2/24	
ge-0/0/0.1	up	up	inet	2.1.1.2/24	
ge-0/0/0.2	up	up	inet	3.1.1.2/24	
ge-0/0/0.3	up	up	inet	4.1.1.2/24	
ge-0/0/0.4	up	up	inet	5.1.1.2/24	
ge-0/0/0.32767	up	up			

Meaning Each subinterface created is displayed as a *ge-fpc/pic/port.x* logical interface, where x is the unit number in the configuration. The status is listed as **up**, indicating the link is working.

Verifying That Traffic Passes Between VLANs

Purpose Verify that the distribution switch is correctly routing traffic from one VLAN to another.

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
```

- 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

- [Example: Connecting an Access Switch to a Distribution Switch](#)
- [Configuring a Layer 3 Subinterface \(CLI Procedure\) on page 105](#)

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 39](#)
- [Overview and Topology on page 40](#)
- [Configuration on page 40](#)
- [Verification on page 41](#)

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 for EX Series Switches” on page 13](#) 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 Procedure To configure unicast RPF on Switch A:

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 41](#)

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:
Total octets      Receive      Transmit
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 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 106](#)
- [Disabling Unicast RPF \(CLI Procedure\) on page 107](#)

Example: Configuring IP Directed Broadcast on an EX Series 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 43](#)
- [Overview and Topology on page 44](#)
- [Configuration on page 44](#)

Requirements

This example uses the following software and hardware components:

- Junos OS Release 9.4 or later for EX Series switches
- One PC
- One EX Series 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 4 on page 44](#)), a host is connected to an interface on an EX Series 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 4: Topology for IP Directed Broadcast



[Table 6 on page 44](#) shows the settings of the components in this example.

Table 6: Components of the IP Directed Broadcast Topology

Property	Settings
Switch hardware	EX Series switch
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 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 vlan.0;
    }
    v1 {
        l3-interface vlan.1;
    }
}
```

Related Documentation • [Configuring IP Directed Broadcast \(CLI Procedure\) on page 108](#)

CHAPTER 3

Configuring Interfaces

- [Configuring Gigabit Ethernet Interfaces \(CLI Procedure\) on page 48](#)
- [Configuring Gigabit Ethernet Interfaces \(J-Web Procedure\) on page 51](#)
- [Port Role Configuration with the J-Web Interface \(with CLI References\) on page 58](#)
- [Adding an Interface Description to the Configuration on page 62](#)
- [Adding a Logical Unit Description to the Configuration on page 63](#)
- [Disabling a Physical Interface on page 64](#)
- [Disabling a Logical Interface on page 65](#)
- [Configuring Flow Control on page 65](#)
- [Configuring the Interface Address on page 66](#)
- [Configuring the Interface Bandwidth on page 68](#)
- [Configuring the Media MTU on page 69](#)
- [Setting the Protocol MTU on page 80](#)
- [Interface Ranges on page 80](#)
- [Configuring Accounting for the Physical Interface on page 90](#)
- [Configuring Accounting for the Logical Interface on page 91](#)
- [Configuring Ethernet Loopback Capability on page 92](#)
- [Configuring Gratuitous ARP on page 93](#)
- [Configuring Static ARP Table Entries on page 94](#)
- [Disabling the Transmission of Redirect Messages on an Interface on page 95](#)
- [Configuring Restricted and Unrestricted Proxy ARP on page 95](#)
- [Enabling or Disabling SNMP Notifications on Logical Interfaces on page 96](#)
- [Enabling or Disabling SNMP Notifications on Physical Interfaces on page 97](#)
- [Configuring Aggregated Ethernet Interfaces \(CLI Procedure\) on page 97](#)
- [Configuring Aggregated Ethernet Interfaces \(J-Web Procedure\) on page 98](#)
- [Configuring Aggregated Ethernet LACP \(CLI Procedure\) on page 101](#)
- [Configuring Aggregated Ethernet Link Protection on page 102](#)
- [Configuring Aggregated Ethernet Link Speed on page 103](#)
- [Configuring Aggregated Ethernet Minimum Links on page 104](#)

- [Configuring Tagged Aggregated Ethernet Interfaces on page 105](#)
- [Configuring a Layer 3 Subinterface \(CLI Procedure\) on page 105](#)
- [Configuring Unicast RPF \(CLI Procedure\) on page 106](#)
- [Disabling Unicast RPF \(CLI Procedure\) on page 107](#)
- [Configuring IP Directed Broadcast \(CLI Procedure\) on page 108](#)
- [Tracing Operations of an Individual Router or Switch Interface on page 109](#)
- [Tracing Operations of the Interface Process on page 110](#)
- [Setting the Mode on an SFP+ Uplink Module \(CLI Procedure\) on page 110](#)
- [Configuring the Media Type on Dual-Purpose Uplink Ports \(CLI Procedure\) on page 111](#)

[Configuring Gigabit Ethernet Interfaces \(CLI Procedure\)](#)

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 48](#)
- [Configuring the Link Settings on page 49](#)
- [Configuring the IP Options on page 51](#)

[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 \(CLI Procedure\)](#).

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 58 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 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 Interfaces \(CLI Procedure\)” on page 97](#).
- **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 51](#)
- [Monitoring Interface Status and Traffic on page 113](#)
- [show interfaces ge- on page 228](#)
- [show interfaces xe- on page 264](#)
- [Understanding Interface Naming Conventions on EX Series Switches on page 6](#)

Configuring Gigabit Ethernet Interfaces (J-Web Procedure)

An Ethernet interface must be configured for optimal performance in a high-traffic network.

To configure properties on a Gigabit Ethernet interface or a 10-Gigabit Ethernet interface on an EX Series switch:

1. Select **Interfaces > Ports**.

The page lists Gigabit Ethernet and 10-Gigabit Ethernet interfaces and their link status.



NOTE: After you make changes to the configuration in 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. Select the interface you want to configure. If the interface you want to configure is not listed under **Ports** in the top table on the page, select the FPC (the FPC is the line card on an EX8200 switch or the member switch in a Virtual Chassis configuration) that includes that interface from the **List Ports for FPC** list.

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



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



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, pre-configured 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 are selecting a port role, ensure that the corresponding port security settings for the VLAN are applicable to the interface.

For basic information on 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 Options—Enables you to configure VLAN options for the selected interface.
 - Link Options—Enables you to modify the following link options for the selected interface:
 - Speed
 - MTU
 - Autonegotiation
 - Flow Control
 - Duplex
 - IP Options—Enables you to configure an IP address for the interface.
4. Configure the interface by configuring options in the selected option set. See [Table 7 on page 54](#) for details on 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 for a selected interface, click **Enable Port** or **Disable Port**.

Table 7: Port Edit Options

Field	Function	Your Action
Port Role	<p>Specifies a profile (role) to assign to the interface.</p> <p>NOTE: Once a port role is configured on the interface, you cannot specify VLAN options or IP options.</p> <p>NOTE: Only the following port roles can be applied on EX8200 switch interfaces:</p> <ul style="list-style-type: none"> • Default • Layer 2 uplink • Routed uplink 	
Default	<p>Applies the default role.</p> <p>The interface family is set to ethernet-switching, port mode is set to access, and RSTP is enabled.</p>	<ol style="list-style-type: none"> 1. Click Details to view CLI commands for this role. 2. Click OK.
Desktop	<p>Applies the desktop role.</p> <p>The interface family is set to ethernet-switching, port mode is set to access, RSTP is enabled with the edge and point-to-point options, and port security parameters (MAC limit =1; dynamic ARP inspection and DHCP snooping enabled) are set.</p>	<ol style="list-style-type: none"> 1. Select an existing VLAN configuration or type the name of a new VLAN configuration to be associated with the interface. 2. Click Details to view CLI commands for this role. 3. Click OK.
Desktop and Phone	<p>Applies the desktop and phone role.</p> <p>The interface family is set to ethernet-switching, port mode is set to access, port security parameters (MAC limit =1; dynamic ARP Inspection and DHCP snooping enabled) are set, and recommended CoS parameters are specified for forwarding classes, schedulers, and classifiers. See Table 8 on page 57 for more CoS information.</p>	<ol style="list-style-type: none"> 1. Select an existing VLAN configuration or type the name of a new VLAN configuration to be associated with the interface. You can also select an existing VoIP VLAN configuration or a new VoIP VLAN configuration to be associated with the interface. NOTE: VoIP is not supported on EX8200 switches. 2. Click Details to view CLI commands for this role. 3. Click OK.
Wireless Access Point	<p>Applies the wireless access point role.</p> <p>The interface family is set to ethernet-switching, port mode is set to access, and RSTP is enabled with the edge and point-to-point options.</p>	<ol style="list-style-type: none"> 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. 2. Click Details to view CLI commands for this role. 3. Click OK.

Table 7: Port Edit Options (*continued*)

Field	Function	Your Action
Routed Uplink	<p>Applies the routed uplink role.</p> <p>The interface family is set to inet, and recommended CoS parameters are set for schedulers and classifiers. See Table 8 on page 57 for more CoS information.</p>	<p>To specify an IPv4 address:</p> <ol style="list-style-type: none"> 1. Select the check box IPv4 address. 2. Type an IP address—for example: 10.10.10.10. 3. Enter the subnet mask or address prefix. For example, 24 bits represents 255.255.255.0. 4. Click OK. <p>To specify an IPv6 address:</p> <ol style="list-style-type: none"> 1. Select the check box IPv6 address. 2. Type an IP address—for example: 2001:ab8:85a3::8a2e:370:7334. 3. Enter the subnet mask or address prefix. 4. Click OK. <p>NOTE: Ipv6 is not supported on EX2200 and EX4500 switches.</p>
Layer 2 Uplink	<p>Applies the Layer 2 uplink role.</p> <p>The interface family is set to ethernet-switching, port mode is set to trunk, RSTP is enabled with the point-to-point option, and port security is set to dhcp-trusted.</p>	<ol style="list-style-type: none"> 1. For this port role you can select a VLAN member and associate a native VLAN with the interface. 2. Click Details to view CLI commands for this role. 3. Click OK.
None	Specifies that no port role is configured for the selected interface.	
<p>NOTE: See “Port Role Configuration with the J-Web Interface (with CLI References)” on page 58 for details on the CLI commands that are associated with each port role.</p> <p>NOTE: For an EX8200 switch, dynamic ARP inspection and DHCP snooping parameters are not configured.</p>		
VLAN Options		

Table 7: 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 Trunk, you can:</p> <ol style="list-style-type: none"> 1. Click Add to add a VLAN member. 2. Select the VLAN and click OK. 3. (Optional) Associate a native VLAN with the interface. 4. Click OK. <p>If you select Access, you can:</p> <ol style="list-style-type: none"> 1. Select the VLAN member to be associated with the interface. 2. (Optional) Associate a VoIP VLAN with the interface. Only a VLAN with a VLAN ID can be associated as a VoIP VLAN. <p>NOTE: VoIP is not supported on EX8200 switches.</p> <ol style="list-style-type: none"> 3. Click OK.
Link Options		
MTU (bytes)	Specifies the maximum transmission unit size for the interface.	Type a value from 256 through 9216 . The default MTU for Gigabit Ethernet interfaces is 1514 .
Speed	Specifies the speed for the mode.	Select one of the following values: 10 Mbps, 100 Mbps, 1000 Mbps, or Auto-Negotiation.
Duplex	Specifies the link mode.	Select one: automatic , half , or full .
Description	<p>Describes the link.</p> <p>NOTE: If the interface is part of a link aggregation group (LAG), only the option Description is enabled.</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.
IP Options		

Table 7: Port Edit Options (*continued*)

Field	Function	Your Action
IPv4 Address	Specifies an IPv4 address for the interface. NOTE: If the IP address is cleared, the interface still belongs to the inet family.	<ol style="list-style-type: none"> To specify an IPv4 address, select the check box IPv4 address. Type an IP address—for example: 10.10.10.10. Enter the subnet mask or address prefix. For example, 24 bits represents 255.255.255.0. Click OK.
IPv6 Address	Specifies an IPv6 address for the interface. NOTE: If the IP address is cleared, the interface still belongs to the inet family.	<ol style="list-style-type: none"> To specify an IPv6 address, select the check box IPv6 address. Type an IP address—for example: 2001:ab8:85a3::8a2e:370:7334. Enter the subnet mask or address prefix. Click OK. <p>NOTE: Ipv6 is not supported on EX2200 and EX4500 switches.</p>

Table 8: 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"> voice—Queue number is set to 7. expedited-forwarding—Queue number is set to 5. assured-forwarding—Queue number is set to 1. best-effort—Queue number is set to 0.
Schedulers	<p>The schedulers and their settings are:</p> <ul style="list-style-type: none"> Strict-priority—Transmission rate is set to 10 percent and buffer size to 5 percent. Expedited-scheduler—Transmission rate is set to 30 percent, buffer size to 30 percent, and priority to low. Assured-scheduler—Transmission rate is set to 25 percent, buffer size to 25 percent, and priority to low. Best-effort scheduler—Transmission rate is set to 35 percent, buffer size to 40 percent, and priority to low.
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 ieee-802.1 classifier configuration and sets the loss priority to low for the code point 101 for the voice forwarding class.
dscp classifier	Imports the default dscp classifier configuration and sets the loss priority to low for the code point 101110 for the voice forwarding class.

Related Documentation

- [Configuring Gigabit Ethernet Interfaces \(CLI Procedure\) on page 48](#)
- [Monitoring Interface Status and Traffic on page 113](#)
- [EX Series Switches Interfaces Overview on page 3](#)
- [Junos OS CoS for EX Series Switches Overview](#)
- [Understanding Interface Naming Conventions on EX Series Switches on page 6](#)

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 9 on page 58](#) 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 9: Port Role Configuration Summary

Configuration Description	CLI Commands
Default Port Role	
Set the port role to Default .	<code>set interfaces <i>interface</i> apply-macro juniper-port-profile Default</code>
Set port family to ethernet-switching .	<code>set interfaces <i>interface</i> unit 0 family ethernet-switching</code>
Set port mode to access .	<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	
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 <<i>vlan name</i>> vlan-id <<i>vlan-id</i>></code>
Set port family to ethernet-switching .	<code>set interfaces <i>interface</i> unit 0 family ethernet-switching</code>
Set Port Mode to Access .	<code>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>

Table 9: Port Role Configuration Summary (*continued*)

Configuration Description	CLI Commands
Set port security parameters.	<code>set ethernet-switching-options secure-access-port vlan MacTest arp-inspection</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>
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.	<code>set vlans <i>vlan-name</i> vlan-id <i>vlan id</i></code>
Set voice VLAN if new voice VLAN is specified.	
Set port family to ethernet-switching .	<code>set interfaces <i>interface</i> unit 0 family ethernet-switching</code>
Set Port Mode to access .	<code>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 SCHEDULER_MAP= juniper-port-profile-map IEEE_CLASSIFIER= juniper-ieee-classifier DSCP_CLASSIFIER= juniper-dscp-classifier	<code>set class-of-service interfaces <i>interface</i> scheduler-map juniper-port-profile-map</code> <code>set class-of-service interfaces <i>interface</i> unit 0 classifiers ieee-802.1 juniper_ieee_classifier</code> <code>set class-of-service interfaces <i>interface</i> unit 0 classifiers dscp juniper-dscp-classifier</code>
Set CoS Configuration	Refer Table 10 on page 61 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>
Set VLAN on VLANs stanza.	<code>set vlans <i>vlan name</i> vlan-id <i>vlan-id</i></code>
Set port family to ethernet-switching	<code>set interfaces <i>interface</i> unit 0 family ethernet-switching</code>
Set port mode to Access .	<code>port-mode access</code>
Set VLAN on port stanza.	<code>set interfaces <i>interface</i> unit 0 family ethernet-switching vlan members <i>vlan-members</i></code>
Set RSTP protocol with edge option.	<code>set protocols rstp interface <i>interface</i> edge</code>

Table 9: Port Role Configuration Summary (*continued*)

Configuration Description	CLI Commands
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 Routed Uplink</code>
Set port family to inet. Set IP address on the port.	<code>set interfaces <i>interface</i> unit 0 family inet address <i>ipaddress</i></code>
Set class-of-service parameters SCHEDULER_MAP= juniper-port-profile-map IEEE_CLASSIFIER= juniper-ieee-classifier DSCP_CLASSIFIER= juniper-dscp-classifier	<code>set class-of-service interfaces <i>interfaces</i> scheduler-map juniper-port-profile-map set class-of-service interfaces <i>interface</i> unit 0 classifiers ieee-802.1 juniper_ieee_classifier set class-of-service interfaces <i>interface</i> unit 0 classifiers dscp juniper-dscp-classifier</code>
Set CoS configuration	Refer Table 10 on page 61 for details.
Layer 2 Uplink Port Role	
Set the port role to Layer 2 Uplink.	<code>set interfaces <i>interface</i> apply-macro juniper-port-profile Layer2 Uplink</code>
Set port family to ethernet-switching Set port mode to trunk .	<code>set interfaces <i>interface</i> unit 0 family ethernet-switching port-mode trunk</code>
Set Native VLAN name.	<code>set interfaces <i>interface</i> unit 0 family ethernet-switching 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 vlan members <i>vlan-members</i></code>
Set port security parameter.	<code>set ethernet-switching-options secure-access-port 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>
Set class-of-service parameters. SCHEDULER_MAP= juniper-port-profile-map IEEE_CLASSIFIER= juniper_ieee_classifier DSCP_CLASSIFIER= juniper_dscp_classifier	<code>set class-of-service interfaces <i>interfaces</i> scheduler-map juniper-port-profile-map set class-of-service interfaces <i>interface</i> unit 0 classifiers ieee-802.1 juniper_ieee_classifier set class-of-service interfaces <i>interface</i> unit 0 classifiers dscp juniper-dscp-classifier</code>
Set CoS configuration	Refer to Table 10 on page 61 for details.

Table 10 on page 61 lists the CLI commands for the recommended CoS settings that are committed when the CoS configuration is set.

Table 10: Recommended CoS Settings for Port Roles

CoS Parameter	CLI Command
Forwarding Classes	
voice	<code>set class-of-service forwarding-classes class voice queue-num 7</code>
expedited-forwarding	<code>set class-of-service forwarding-classes class expedited-forwarding queue-num 5</code>
assured-forwarding	<code>set class-of-service forwarding-classes class assured-forwarding queue-num 1</code>
best-effort	<code>set class-of-service forwarding-classes class best-effort queue-num 0</code>
Schedulers	
strict-priority-scheduler	<p>The CLI commands are:</p> <ul style="list-style-type: none"> <code>set class-of-service schedulers strict-priority-scheduler transmit-rate percent 10</code> <code>set class-of-service schedulers strict-priority-scheduler buffer-size percent 5</code> <code>set class-of-service schedulers strict-priority-scheduler priority strict-high</code>
expedited-scheduler	<p>The CLI commands are:</p> <ul style="list-style-type: none"> <code>set class-of-service schedulers expedited-scheduler transmit-rate percent 30</code> <code>set class-of-service schedulers expedited-scheduler buffer-size percent 30</code> <code>set class-of-service schedulers expedited-scheduler priority low</code>
assured-scheduler	<p>The CLI commands are:</p> <pre>set class-of-service schedulers assured-scheduler transmit-rate percent 25 set class-of-service schedulers strict-priority-scheduler buffer-size percent 25 set class-of-service schedulers strict-priority-scheduler priority low</pre>
best-effort-scheduler	<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>
Classifiers	<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 51](#)
 - [Configuring Gigabit Ethernet Interfaces \(CLI Procedure\) on page 48](#)

Adding an Interface Description to the Configuration

You can include a text description of each physical interface 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 at the **[edit interfaces *interface-name*]** hierarchy level:

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

```
description text;
```

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 [Enabling and Disabling Insertion of Option 82 Information in the Junos OS Subscriber Access Configuration Guide](#).

For information about describing logical units, see [“Adding a Logical Unit Description to the Configuration” on page 63](#).

Example: Adding an Interface Description to the Configuration

Add a description to a Fast Ethernet interface:

```
[edit interfaces]
user@host#

set fe-0/0/1 description "Backbone connection to PHL01"
[edit interfaces]
user@host#

show
fe-0/0/1 {
  description "Backbone connection to PHL01";
  unit 0 {
    family inet {
      address 192.168.0.1/30;
    }
  }
}
```

To display the description from the router or switch CLI, use the **show interfaces** command:

```
user@host>

show interfaces fe-0/0/1
Physical interface: fe-0/0/1, Enabled, Physical link is Up
  Interface index: 129, SNMP ifIndex: 23
  Description: Backbone connection to PHL01
  ...
```

To display the interface description from the interfaces MIB, use the **snmpwalk** command from a server. To isolate information for a specific interface, search for the interface index

shown in the **SNMP ifIndex** field of the **show interfaces** command output. The **ifAlias** object is in **ifXTable**.

```
user-server>snmpwalk host-fxp0.mylab public ifXTable | grep -e '\.23'
snmpwalk host-fxp0.mylab public ifXTable | grep -e '\.23'
ifMIB.ifMIBObjects.ifXTable.ifXEntry.ifName.23 = fe-0/0/1
ifMIB.ifMIBObjects.ifXTable.ifXEntry.ifInMulticastPkts.23 = Counter32: 0
ifMIB.ifMIBObjects.ifXTable.ifXEntry.ifInBroadcastPkts.23 = Counter32: 0
ifMIB.ifMIBObjects.ifXTable.ifXEntry.ifOutMulticastPkts.23 = Counter32: 0
ifMIB.ifMIBObjects.ifXTable.ifXEntry.ifOutBroadcastPkts.23 = Counter32: 0
ifMIB.ifMIBObjects.ifXTable.ifXEntry.ifHCInOctets.23 = Counter64: 0
ifMIB.ifMIBObjects.ifXTable.ifXEntry.ifHCInUcastPkts.23 = Counter64: 0
ifMIB.ifMIBObjects.ifXTable.ifXEntry.ifHCInMulticastPkts.23 = Counter64: 0
ifMIB.ifMIBObjects.ifXTable.ifXEntry.ifHCInBroadcastPkts.23 = Counter64: 0
ifMIB.ifMIBObjects.ifXTable.ifXEntry.ifHCOutOctets.23 = Counter64: 42
ifMIB.ifMIBObjects.ifXTable.ifXEntry.ifHCOutUcastPkts.23 = Counter64: 0
ifMIB.ifMIBObjects.ifXTable.ifXEntry.ifHCOutMulticastPkts.23 = Counter64: 0
ifMIB.ifMIBObjects.ifXTable.ifXEntry.ifHCOutBroadcastPkts.23 = Counter64: 0
ifMIB.ifMIBObjects.ifXTable.ifXEntry.ifLinkUpDownTrapEnable.23 = enabled(1)
ifMIB.ifMIBObjects.ifXTable.ifXEntry.ifHighSpeed.23 = Gauge32: 100
ifMIB.ifMIBObjects.ifXTable.ifXEntry.ifPromiscuousMode.23 = false(2)
ifMIB.ifMIBObjects.ifXTable.ifXEntry.ifConnectorPresent.23 = true(1)
ifMIB.ifMIBObjects.ifXTable.ifXEntry.ifAlias.23 = Backbone connection to PHL01
ifMIB.ifMIBObjects.ifXTable.ifXEntry.ifCounterDiscontinuityTime.23 = Timeticks:
(0) 0:00:00.00
```

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 “Enabling and Disabling Insertion of Option 82 Information” in the *Junos OS Subscriber Access Configuration Guide*.

For information about describing physical interfaces, see “Adding an Interface Description to the Configuration” on page 62.

Disabling a Physical Interface

You can disable a physical interface, marking it as being down, without removing the interface configuration statements from the configuration. To do this, include the **disable** statement at the **[edit interfaces *interface-name*]** hierarchy level:

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



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.



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

Disable a physical interface:

```
[edit interfaces]  
so-1/1/0 {  
  mtu 8000;  
  clocking internal;  
  encapsulation ppp;  
  sonet-options {  
    fcs 16;  
  }  
  unit 0 {  
    family inet {  
      address 172.16.0.0/12 {  
        destination 172.16.0.4;  
      }  
    }  
  }  
}  
[edit interfaces]  
user@host# set so-1/1/0 disable  
[edit interfaces]
```

```

user@host# show so-1/1/0
so-1/1/0 {
  disable;# Interface is marked as disabled
  mtu 8000;
  clocking internal;
  encapsulation ppp;
  sonet-options {
    fcs 16;
  }
  unit 0 {
    family inet {
      address 172.16.0.0 {
        destination 172.16.0.3;
      }
    }
  }
}

```

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* gigheter-options]

Related
Documentation

- [flow-control on page 157](#)
- [Ethernet Interfaces Overview](#)
- [Junos OS Ethernet Interfaces Configuration Guide](#)

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 **vpls** 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, include the **address** statement:

```
address address {  
  broadcast address;  
  destination address;  
  destination-profile name;  
  eui-64;  
  preferred;  
  primary;  
}
```

You can include these statements 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*]

In the **address** statement, specify the network address of the interface.

For each address, you can optionally configure one or more of the following:

- Broadcast address for the interface subnet—Specify this in the **broadcast** statement; this applies only to Ethernet interfaces, such as the management interface **fxp0**, **em0**, or **me0** the Fast Ethernet interface, and the Gigabit Ethernet interface.
- Address of the remote side of the connection (for point-to-point interfaces only)—Specify this in the **destination** statement.
- PPP properties to the remote end—Specify this in the **destination-profile** statement. You define the profile at the **[edit access group-profile name ppp]** hierarchy level (for point-to-point interfaces only).
- Whether the router or switch automatically generates the host number portion of interface addresses—The **eui-64** statement applies only to interfaces that carry IPv6 traffic, in which the prefix length of the address is 64 bits or less, and the low-order 64 bits of the address are zero. This option does not apply to the loopback interface (**lo0**) because IPv6 addresses configured on the loopback interface must have a 128-bit prefix length.



NOTE: IPv6 is not currently supported for the QFX Series.

- Whether this address is the preferred address—Each subnet on an interface has a preferred local address. If you configure more than one address on the same subnet, the preferred local address is chosen by default as the source address when you originate packets to destinations on the subnet.

By default, the preferred address is the lowest-numbered address on the subnet. To override the default and explicitly configure the preferred address, include the **preferred** statement when configuring the address.

- Whether this address is the primary address—Each interface has a primary local address. If an interface has more than one address, the primary local address is used by default as the source address when you send packets from an interface where the destination provides no information about the subnet (for example, some **ping** commands).

By default, the primary address on an interface is the lowest-numbered non-127 (in other words, non-loopback) preferred address on the interface. To override the default and explicitly configure the preferred address, include the **primary** statement when configuring the address.

- [Configuring Interface IPv4 Addresses on page 67](#)
- [Configuring Interface IPv6 Addresses on page 68](#)

Configuring Interface IPv4 Addresses

You can configure router or switch interfaces with a 32-bit IP version 4 (IPv4) address and optionally with a destination prefix, sometimes called a *subnet mask*. An IPv4 address utilizes a 4-octet dotted decimal address syntax (for example, **192.16.1.1**). An IPv4 address with destination prefix utilizes a 4-octet dotted decimal address syntax with a destination prefix appended (for example, **192.16.1.1/30**).

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



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.

Configuring Interface IPv6 Addresses

You represent IP version 6 (IPv6) addresses in hexadecimal notation using a colon-separated list of 16-bit values.

You assign a 128-bit IPv6 address to an interface by including the **address** statement:

```
address aaaa:bbbb:...:zzzz/nn;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family inet6]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet6]

The double colon (::) represents all bits set to 0, as shown in the following example:

```
interfaces fe-0/0/1 {  
  unit 0 {  
    family inet6 {  
      address fec0:1:1::2/64;  
    }  
  }  
}
```



NOTE: You must manually configure the router or switch advertisement and advertise the default prefix for autoconfiguration to work on a specific interface.

Related Documentation

- Configuring IPCP Options
- 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 default media MTU size used on a physical interface depends on the encapsulation used on that interface. In some cases, the default IP Protocol MTU depends on whether the protocol used is IP version 4 (IPv4) or International Organization for Standardization (ISO). [Table 11 on page 70](#) through [Table 21 on page 77](#) list the media and protocol MTU sizes by interface type, and [Table 22 on page 77](#) lists the encapsulation overhead by encapsulation type.

Table 11: Media MTU Sizes by Interface Type for M5, M7i with CFEB, M10, M10i with CFEB, M20, and M40 Routers

Interface Type	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP Protocol MTU (Bytes)
Adaptive Services (MTU size not configurable)	9192	N/A	N/A
ATM	4482	9192	4470
E1/T1	1504	9192	1500
E3/T3	4474	9192	4470
Fast Ethernet	1514	9192 (4-port) 1532 (8-port) 1532 (12-port)	1500 (IPv4), 1497 (ISO)
Gigabit Ethernet	1514	9192	1500 (IPv4), 1497 (ISO)
Serial	1504	9192	1500 (IPv4), 1497 (ISO)
SONET/SDH	4474	9192	4470

Table 12: Media MTU Sizes by Interface Type for M40e Routers

Interface Type	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP Protocol MTU (Bytes)
Adaptive Services (MTU size not configurable)	9192	N/A	N/A
ATM	4482	9192	4470
E1/T1	1504	4500	1500
E3/T3	4474	4500 9192 (4-port)	4470
E3/DS3 IQ	4474	9192	4470
Fast Ethernet	1514	4500	1500 (IPv4), 1497 (ISO)

Table 12: Media MTU Sizes by Interface Type for M40e Routers (*continued*)

Interface Type	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP Protocol MTU (Bytes)
Gigabit Ethernet	1514	9192 (1- or 2-port) 9192 (4-port)	1500 (IPv4), 1497 (ISO)
Serial	1504	9192	1500 (IPv4), 1497 (ISO)
SONET/SDH	4474	4500 (1-port nonconcatenated) 9192 (4-port OC3) 9192 (4-port OC3c) 4500 (1-port OC12) 4500 (4-port OC12) 4500 (4-port OC12c) 4500 (1-port OC48) 9192 (2-port OC3) 9192 (2-port OC3c) 9192 (1-port OC12c) 9192 (1-port OC48c) 4500 (1-port OC192) 9192 (1-port OC192c)	4470

Table 13: Media MTU Sizes by Interface Type for M160 Routers

Interface Type	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP Protocol MTU (Bytes)
Adaptive Services (MTU size not configurable)	9192	N/A	N/A
ATM	4482	9192	4470
E1/T1	1504	4500	1500
E3/T3	4474	4500	4470
E3/DS3 IQ	4474	9192	4470
Fast Ethernet	1514	4500	1500 (IPv4), 1497 (ISO)

Table 13: Media MTU Sizes by Interface Type for M160 Routers (*continued*)

Interface Type	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP Protocol MTU (Bytes)
Gigabit Ethernet	1514	9192 (1- or 2-port) 4500 (4-port)	1500 (IPv4), 1497 (ISO)
Serial	1504	9192	1500 (IPv4), 1497 (ISO)
SONET/SDH	4474	4500 (1-port nonconcatenated) 9192 (1- or 2-port) 4500 (4-port)	4470

Table 14: Media MTU Sizes by Interface Type for M7i with CFEB-E, M10i with CFEB-E, M320 and M120 Routers

Interface Type	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP Protocol MTU (Bytes)
ATM2 IQ	4482	9192	4470
Channelized DS3 IQ	4471	4500	4470
Channelized E1 IQ	1504	4500	1500
Channelized OC12 IQ	4474	9192	4470
Channelized STM1 IQ	4474	9192	4470
DS3	4471	4500	4470
E1	1504	4500	1500
E3 IQ	4471	4500	4470
Fast Ethernet	1514	9192 (4-port) 1532 (8-, 12- and 48-port)	1500 (IPv4), 1497 (ISO)
Gigabit Ethernet	1514	9192	1500 (IPv4), 1497 (ISO)
SONET/SDH	4474	9192	4470
T1	1504	4500	1500

Table 14: Media MTU Sizes by Interface Type for M7i with CFEB-E, M10i with CFEB-E, M320 and M120 Routers (*continued*)

Interface Type	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP Protocol MTU (Bytes)
CT3 IQ (excluding M120)	4474	9192	4470

Table 15: Media MTU Sizes by Interface Type for MX Series Routers

Interface Type	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP Protocol MTU (Bytes)
Gigabit Ethernet	1514	9192	1500 (IPv4) 1488 (MPLS) 1497 (ISO)
10-Gigabit Ethernet	1514	9192	1500 (IPv4) 1488 (MPLS) 1497 (ISO)
Multi-Rate Ethernet	1514	9192	1500 (IPv4) 1488 (MPLS) 1497 (ISO)
Tri-Rate Ethernet	1514	9192	1500 (IPv4) 1488 (MPLS) 1497 (ISO)
Channelized SONET/SDH OC3/STM1 (Multi-Rate)	1514	9192	1500 (IPv4) 1488 (MPLS) 1497 (ISO)
DS3/E3 (Multi-Rate)	1514	9192	1500 (IPv4) 1488 (MPLS) 1497 (ISO)

Table 16: Media MTU Sizes by Interface Type for T320 Routers

Interface Type	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP Protocol MTU (Bytes)
ATM	4482	9192	4470
ATM2 IQ	4482	9192	4470
Channelized OC12 IQ	4474	9192	4470
Channelized STM1 IQ	4474	9192	4470
DS3	4471	4500	4470

Table 16: Media MTU Sizes by Interface Type for T320 Routers (*continued*)

Interface Type	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP Protocol MTU (Bytes)
Fast Ethernet	1514	4500 (4-port) 1532 (12- and 48-port)	1500 (IPv4), 1497 (ISO)
Gigabit Ethernet	1514	9192	1500 (IPv4), 1497 (ISO)
SONET/SDH	4474	9192	4470
CT3 IQ	4474	9192	4470

Table 17: Media MTU Sizes by Interface Type for T640 Platforms

Interface Type	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP Protocol MTU (Bytes)
ATM2 IQ	4482	9192	4470
48-port Fast Ethernet	1514	1532	1500 (IPv4), 1497 (ISO)
Gigabit Ethernet	1514	9192	1500 (IPv4), 1497 (ISO)
SONET/SDH	4474	9192	4470
CT3 IQ	4474	9192	4470

Table 18: Media MTU Sizes by Interface Type for J2300 Platforms

Interface Type	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP Protocol MTU (Bytes)
Fast Ethernet (10/100)	1514	9192	1500
G.SHDSL	4482	9150	4470
ISDN BRI	1504	4092	1500
Serial	1504	9150	1500
T1 or E1	1504	9150	1500

Table 19: Media MTU Sizes by Interface Type for J4300 and J6300 Platforms

Interface Type	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP Protocol MTU (Bytes)
ADSL2+ PIM	4482	9150	4470
Dual-port Fast Ethernet (10/100) PIM	1514	9192	1500
Dual-port Serial PIM	1504	9150	1500
Dual-port T1 or E1 PIM	1504	9150	1500
Dual-port Channelized T1/E1 PIM (channelized to DS0s)	1504	4500	1500
Dual-port Channelized T1/E1 PIM (clear channel T1 or E1)	1504	9150	1500
Fast Ethernet (10/100) built-in interface	1514	9192	1500
G.SHDSL PIM	4482	9150	4470
4-port ISDN BRI PIM	1504	4092	1500
T3 (DS3) or E3 PIM	4474	9192	4470

Table 20: Media MTU Sizes by Interface Type for J4350 and J6350 Platforms

Interface Type	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP Protocol MTU (Bytes)
4-port ISDN BRI PIM	1504	4092	1500
ADSL2+ PIM	4482	9150	4470
Dual-port Fast Ethernet (10/100) PIM	1514	9192	1500
Dual-port Serial PIM	1504	9150	1500
Dual-port T1 or E1 PIM	1504	9150	1500

Table 20: Media MTU Sizes by Interface Type for J4350 and J6350 Platforms (*continued*)

Interface Type	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP Protocol MTU (Bytes)
Dual-port Channelized T1/E1 PIM (channelized to DS0s)	1504	4500	1500
Dual-port Channelized T1/E1 PIM (clear channel T1 or E1)	1504	9150	1500
4-port Fast Ethernet (10/100) ePIM	1518	1518	1500
Gigabit Ethernet (10/100/1000) built-in interface	1514	9018	1500
Gigabit Ethernet (10/100/1000) Enhanced Physical Interface Module (ePIM)	1514	9018	1500
Gigabit Ethernet (10/100/1000) SFP ePIM	1514	9018	1500
G.SHDSL PIM	4482	9150	4470
T3 (DS3) or E3 PIM	4474	9192	4470



NOTE: On Gigabit Ethernet ePIMs in J4350 and J6350 Services Routers, you can configure a maximum transmission unit (MTU) size of only 9018 bytes even though the CLI indicates that you can configure an MTU of up to 9192 bytes. If you configure an MTU greater than 9018 bytes, the router does not accept the configuration and generates a system log error message similar to the following:

```
/kernel: ge-0/0/0: Illegal media change. MTU invalid: 9192. Max MTU supported on this PIC: 9018
```

On 4-port Fast Ethernet ePIMs in J4350 and J6350 Services Routers, you can configure a maximum transmission unit (MTU) size of only 1518 bytes even though the CLI indicates that you can configure an MTU of up to 9192 bytes. If you configure an MTU greater than 1518 bytes, the router does not accept the configuration and generates a system log error message similar to the following:

```
/kernel: fe-3/0/1: Illegal media change. MTU invalid: 9192. Max MTU supported on this PIC: 1518
```

Table 21: Media MTU Sizes by Interface Type for EX Series Switches

Interface Type	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP Protocol MTU (Bytes)
Gigabit Ethernet	1514	9192	1500 (IPv4), 1497 (ISO)
10-Gigabit Ethernet	1514	9192	1500 (IPv4), 1497 (ISO)

Table 22: Encapsulation Overhead by Encapsulation Type

Interface Encapsulation	Encapsulation Overhead (Bytes)
802.1Q/Ethernet 802.3	21
802.1Q/Ethernet Subnetwork Access Protocol (SNAP)	26
802.1Q/Ethernet version 2	18
ATM Cell Relay	4
ATM permanent virtual connection (PVC)	12
Cisco HDLC	4
Ethernet 802.3	17
Ethernet circuit cross-connect (CCC) and virtual private LAN service (VPLS)	4

Table 22: Encapsulation Overhead by Encapsulation Type (*continued*)

Interface Encapsulation	Encapsulation Overhead (Bytes)
Ethernet over ATM	32
Ethernet SNAP	22
Ethernet translational cross-connect (TCC)	18
Ethernet version 2	14
Extended virtual local area network (VLAN) CCC and VPLS	4
Extended VLAN TCC	22
Frame Relay	4
PPP	4
VLAN CCC	4
VLAN VPLS	4
VLAN TCC	22

The default media MTU is calculated as follows:

Default media MTU = Default IP MTU + encapsulation overhead

When you are configuring point-to-point connections, the MTU sizes on both sides of the connections must be the same. Also, when you are configuring point-to-multipoint connections, all interfaces in the subnet must use the same MTU size.



NOTE: The actual frames transmitted also contain cyclic redundancy check (CRC) bits, which are not part of the media MTU. For example, the media MTU for a Gigabit Ethernet Version 2 interface is specified as 1514 bytes, but the largest possible frame size is actually 1518 bytes; you need to consider the extra bits in calculations of MTUs for interoperability.

The physical MTU for Ethernet interfaces does not include the 4-byte frame check sequence (FCS) field of the Ethernet frame.

A SONET/SDH interface operating in concatenated mode has a “c” added to the rate descriptor. For example, a concatenated OC48 interface is referred to as OC48c.

If you do not configure an MPLS MTU, the Junos OS derives the MPLS MTU from the physical interface MTU. From this value, the software subtracts the encapsulation-specific overhead and space for the maximum number of labels that might be pushed in the Packet Forwarding Engine. Currently, the software provides for three labels of four bytes each, for a total of 12 bytes.

In other words, the formula used to determine the MPLS MTU is the following:

$$\text{MPLS MTU} = \text{physical interface MTU} - \text{encapsulation overhead} - 12$$

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.

For information about configuring the encapsulation on an interface, see [Configuring Interface Encapsulation on Physical Interfaces](#).

To modify the default media MTU size for a physical interface, include the `mtu` statement at the `[edit interfaces interface-name]` hierarchy level:

```
[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.



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

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]`

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, as described in “[Setting the Protocol MTU](#)” on page 80.

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 “[Configuring the Media MTU](#)” on page 69.

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 [Table 22 on page 77](#). 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, as discussed in “[Configuring the Media MTU](#)” on page 69.)



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.

Interface Ranges

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 81](#)
- [Expanding Interface Range Member and Member Range Statements on page 84](#)
- [Configuration Inheritance for Member Interfaces on page 86](#)
- [Member Interfaces Inheriting Configuration from Configuration Groups on page 87](#)
- [Interfaces Inheriting Common Configuration on page 88](#)
- [Configuring Inheritance Range Priorities on page 88](#)
- [Configuration Expansion Where Interface Range Is Used on page 89](#)

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 all available interfaces.
- **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 [Junos OS 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**.

```
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 the 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 below:

```

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.

Configuring Accounting for the Physical Interface

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 [Junos OS Network Management Configuration Guide](#).

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 [Junos OS Routing Policy Configuration Guide](#).

Applying an Accounting Profile to the Physical Interface

To enable accounting on an interface, include the **accounting-profile** statement at the **[edit interfaces *interface-name*]** hierarchy level:

```
[edit interfaces interface-name]  
  accounting-profile name;
```

You can also reference profiles by logical unit; for more information, see “[Configuring Accounting for the Logical Interface](#)” on page 91.

Example: Applying an Accounting Profile to the Physical Interface

Configure an accounting profile for an interface and apply it to a physical interface:

```
[edit]  
accounting-options {  
  file if_stats {  
    size 4m files 10 transfer-interval 15;  
    archive-sites {  
      "ftp://login:password@host/path";  
    }  
  }  
  interface-profile if_profile {  
    interval 15;  
    file if_stats {  
      fields {  
        input-bytes;  
        output-bytes;  
      }  
    }  
  }  
}
```



```

        input-packets;
        output-packets;
        input-errors;
        output-errors;
    }
}
}
[edit interfaces ge-1/0/1]
accounting-profile if_profile;

```

Configuring Accounting for the Logical Interface

Juniper Networks routers or switches can collect various kinds of data about traffic passing through the router or 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 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 [Junos OS Network Management Configuration Guide](#).

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 [Junos OS Routing Policy Configuration Guide](#).

Applying an Accounting Profile to the Logical Interface

To enable accounting on a logical interface, include the **accounting-profile** statement:

```
accounting-profile name;
```

You can include this statement at the following hierarchy level:

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

You can also reference profiles for the physical interface; for more information, see “Configuring Accounting for the Physical Interface” on page 90.

Example: Applying an Accounting Profile to the Logical Interface

Configure an accounting profile for an interface and apply it to a logical interface:

```

[edit]
accounting-options {

```

```
file if_stats {
  size 4m files 10 transfer-interval 15;
  archive-sites {
    "ftp://login:password@host/path";
  }
}
interface-profile if_profile {
  interval 15;
  file if_stats {
    fields {
      input-bytes;
      output-bytes;
      input-packets;
      output-packets;
      input-errors;
      output-errors;
    }
  }
}
[edit interfaces ge-1/0/1 unit 1]
accounting-profile if_profile;
```

To reference profiles by physical interface, see [“Applying an Accounting Profile to the Physical Interface” on page 90](#). For information about configuring a firewall filter accounting profile, see the *Junos OS Routing Policy Configuration Guide*.

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/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 173](#)
- Ethernet Interfaces Overview
- [Junos OS Ethernet Interfaces Configuration Guide](#)

Configuring Gratuitous ARP

Gratuitous Address Resolution Protocol (ARP) requests provide duplicate IP address detection. A gratuitous ARP request 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. 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.

By default, the router or switch responds to gratuitous ARP requests. On Ethernet interfaces, you can disable responses to gratuitous ARP requests. To disable responses to gratuitous ARP requests, include the **no-gratuitous-arp-request** statement at the [edit interfaces *interface-name*] hierarchy level:

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

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. On Ethernet interfaces, you can enable handling of gratuitous ARP replies on a specific interface by including the **gratuitous-arp-reply** statement at the [edit interfaces *interface-name*] hierarchy level:

```
[edit interfaces interface-name]
gratuitous-arp-reply;
```

To restore the default behavior, include the **no-gratuitous-arp-reply** statement at the **[edit interfaces *interface-name*]** hierarchy level:

```
[edit interfaces interface-name]  
no-gratuitous-arp-reply;
```

**Related
Documentation**

- [gratuitous-arp-reply on page 158](#)
- no-gratuitous-arp-request
- Ethernet Interfaces Overview
- [Junos OS Ethernet Interfaces Configuration Guide](#)

Configuring Static ARP Table Entries

To configure static ARP table entries, include the **arp** statement:

```
arp ip-address (mac | multicast-mac) mac-address <publish>;
```

You can include this statement at the following hierarchy levels:

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

The IP address that you specify must be part of the subnet defined in the enclosing **address** statement.

To associate a multicast MAC address with a unicast IP address, include the **multicast-mac** statement.

Specify the MAC address as six hexadecimal bytes in one of the following formats: *nnnn.nnnn.nnnn* or *nn:nn:nn:nn:nn:nn*; for example, 0011.2233.4455 or 00:11:22:33:44:55.

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



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.

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.



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

Example: Configuring Static ARP Table Entries

Configure two static ARP table entries on the router or switch's management interface:

```
[edit interfaces]
fxp0 {
  unit 0 {
    family inet {
      address 10.10.0.11/24 {
        arp 10.10.0.99 mac 0001.0002.0003;
        arp 10.10.0.101 mac 00:11:22:33:44:55 publish;
      }
    }
  }
}
```

Related Documentation

- Management Ethernet Interface Overview
- Applying Policers
- Configuring an Unnumbered Interface
- [Junos OS Ethernet Interfaces Configuration Guide](#)

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.

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 responds to any ARP request as long as the router 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 interface with the **no-gratuitous-arp-reply** statement. See “Configuring Gratuitous ARP” on [page 93](#) for information about how to disable responses to gratuitous ARP requests.

Related Documentation

- [proxy-arp on page 187](#)
- Restricted and Unrestricted Proxy ARP Overview
- [Configuring Gratuitous ARP on page 93](#)
- [Junos OS Ethernet Interfaces Configuration Guide](#)

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*]



NOTE: Gigabit Ethernet interfaces on J Series routers do not support SNMP.

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. To explicitly enable these notifications on the physical interface, include the **traps** statement at the **[edit interfaces *interface-name*]** hierarchy level. To disable these notifications on the physical interface, include the **no-traps** statement at the **[edit interfaces *interface-name*]** hierarchy level:

```
[edit interfaces interface-name]  
(traps | no-traps);
```



NOTE: Gigabit Ethernet interfaces on J Series routers do not support SNMP.

Configuring Aggregated Ethernet Interfaces (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. Link aggregation increases bandwidth, provides graceful degradation as failure occurs, and increases 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 2
```

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 2
```

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

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

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

```
[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
```

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

```
[edit interfaces]  
user@switch# set ae0 unit 0 family inet address 192.0.2.0/25
```

For information about adding LACP to a LAG, see “[Configuring Aggregated Ethernet LACP \(CLI Procedure\)](#)” on page 101.

Related Documentation

- [Configuring Aggregated Ethernet Interfaces \(J-Web Procedure\)](#) on page 98
- [Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch](#) on page 21
- [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 27
- [Verifying the Status of a LAG Interface](#) on page 114
- [Understanding Aggregated Ethernet Interfaces and LACP](#) on page 8

Configuring Aggregated Ethernet Interfaces (J-Web Procedure)

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 in 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 23 on page 99](#).
- **Edit**—Modifies a selected LAG.
 - **Aggregation**—Modifies settings for the selected LAG. Enter information as specified in [Table 23 on page 99](#).
 - **VLAN**—Specifies VLAN options for the selected LAG. Enter information as specified in [Table 24 on page 100](#).
 - **IP Option**—Specifies IP options for the selected LAG. Enter information as specified in [Table 25 on page 100](#).
- **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 23: 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"> • None—Indicates that no mode is applicable. • Active—Indicates that the interface initiates transmission of LACP packets • Passive—Indicates that the interface responds only to LACP packets. 	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 Add. Click OK.</p> <p>To remove an interface from the LAG, select the interface and click Remove.</p> <p>NOTE: 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 24: VLAN Options

Field	Function	Your Action
Port Mode	Specifies the mode of operation for the port: trunk or access.	<p>If you select Trunk, you can:</p> <ol style="list-style-type: none"> 1. Click Add to add a VLAN member. 2. Select the VLAN and click OK. 3. (Optional) Associate a native VLAN ID with the port. <p>If you select Access, you can:</p> <ol style="list-style-type: none"> 1. Select the VLAN member to be associated with the port. 2. (Optional) Associate a VoIP VLAN with the interface. Only a VLAN with a VLAN ID can be associated as a VoIP VLAN. <p>Click OK.</p>

Table 25: IP Options

Field	Function	Your Action
IPv4 Address	Specifies an IPv4 address for the selected LAG.	<ol style="list-style-type: none"> 1. Select the check box IPv4 address. 2. Type an IP address—for example, 10.10.10.10. 3. Enter the subnet mask or address prefix. For example, 24 bits represents 255.255.255.0. 4. Click OK.
IPv6 Address	Specifies an IPv6 address for the selected LAG.	<ol style="list-style-type: none"> 1. Select the check box IPv6 address. 2. Type an IP address—for example, 2001:ab8:85a3::8a2e:370:7334. 3. Enter the subnet mask or address prefix. 4. Click OK.

- Related Documentation**
- [Configuring Aggregated Ethernet Interfaces \(CLI Procedure\) on page 97](#)
 - [Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch on page 21](#)
 - [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 27](#)
 - [Verifying the Status of a LAG Interface on page 114](#)

- [Configuring Aggregated Ethernet LACP \(CLI Procedure\) on page 101](#)
- [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 with or without LACP enabled.

Before you configure LACP, be sure you have:

- Configured the aggregated Ethernet bundles—also known as link aggregation groups (LAGs). See [“Configuring Aggregated Ethernet Interfaces \(CLI Procedure\)” on page 97](#)

When LACP is enabled, the local and remote sides of the aggregated Ethernet links exchange protocol data units (PDUs), containing 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. Enable 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 fast
```



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 Interfaces \(CLI Procedure\) on page 97](#)
- [Configuring Aggregated Ethernet Interfaces \(J-Web Procedure\) on page 98](#)
- [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 27](#)

- [Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch on page 21](#)
- [Verifying the Status of a LAG Interface on page 114](#)
- [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.

- [Configuring Link Protection for Aggregated Ethernet Interfaces on page 102](#)
- [Configuring Primary and Backup Links for Link Aggregated Ethernet Interfaces on page 102](#)
- [Reverting Traffic to a Primary Link When Traffic is Passing Through a Backup Link on page 103](#)
- [Disabling Link Protection for Aggregated Ethernet Interfaces on page 103](#)

Configuring Link Protection for Aggregated Ethernet Interfaces

Aggregated Ethernet interfaces support link protection to ensure QoS on the interface.

To configure link protection:

1. Specify that you want to configure the options for an aggregated Ethernet interface.

```
user@host# edit interfaces aex aggregated-ether-options
```

2. Configure the link protection mode.

```
[edit interfaces aex aggregated-ether-options]  
user@host# set link-protection
```

Configuring Primary and Backup Links for Link Aggregated Ethernet Interfaces

To configure link protection, you must specify a primary and a secondary, or backup, link.

To configure a primary link and a backup link:

1. Configure the primary logical interface.

```
[edit interfaces interface-name]  
user@host# set (fastether-options | gigether-options) 802.3ad aex primary
```

2. Configure the backup logical interface.

```
[edit interfaces interface-name]
user@host# set (fastether-options | gigether-options) 802.3ad aex backup
```

Reverting Traffic to a Primary Link When Traffic is Passing Through a Backup Link

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.

To manually control when traffic should be diverted back to the primary link from the designated backup link, enter the following operational command:

```
user@host> request interface revert aex
```

Disabling Link Protection for Aggregated Ethernet Interfaces

To disable link protection, issue the **delete interface revert aex** configuration command.

```
user@host# delete interfaces aex aggregated-ether-options link-protection
```

Configuring Aggregated Ethernet Link Speed

On aggregated Ethernet interfaces, you can set the required link speed for all interfaces included in the bundle. All interfaces that make up a bundle must be the same speed. If you include in the aggregated Ethernet interface an individual link that has a speed different from the speed you specify in the **link-speed** parameter, an error message will be logged.

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).

Aggregated Ethernet interfaces on the M120 router can have one of the following speed values:

- **100m**—Links are 100 Mbps.
- **10g**—Links are 10 Gbps.
- **1g**—Links are 1 Gbps.
- **OC192**—Links are OC192 or STM64c.

Aggregated Ethernet links on EX Series switches can be configured to operate at one of the following speeds:

- **10m**
- **100m**
- **1g**
- **10g**
- **50g**

**Related
Documentation**

- aggregated-ether-options
- [Junos OS Ethernet Interfaces Configuration Guide](#)

Configuring Aggregated Ethernet Minimum Links

On aggregated Ethernet interfaces, you can configure the minimum number of links that must be up for the bundle as a whole to be labeled **up**. By default, only one link must be up for the bundle to be labeled **up**.

To configure the minimum number of links:

1. Specify that you want to configure the aggregated Ethernet options.

```
user@host# edit interfaces interface-name aggregated-ether-options
```

2. Configure the minimum number of links.

```
[edit interfaces interface-name aggregated-ether-options]  
user@host# set minimum-links number
```

On M120, M320, MX Series, T Series, and TX Matrix routers with Ethernet interfaces, the valid range for **minimum-links** *number* is 1 through 16. When the maximum value (16) is specified, all configured links of a bundle must be up for the bundle to be labeled **up**.

On all other routers and on EX Series switches, other than EX8200 switches, the range of valid values for **minimum-links** *number* is 1 through 8. When the maximum value (8) is specified, all configured links of a bundle must be up for the bundle to be labeled **up**.

On EX8200 switches, the range of valid values for **minimum-links** *number* is 1 through 12. When the maximum value (12) is specified, all configured links of a bundle must be up for the bundle to be labeled **up**.

If the number of links configured in an aggregated Ethernet interface is less than the minimum link value configured under the **aggregated-ether-options** statement, the configuration commit fails and an error message is displayed.

- Related Documentation**
- [aggregated-ether-options](#)
 - [minimum-links on page 178](#)
 - [Junos OS Ethernet Interfaces Configuration Guide](#)

Configuring Tagged Aggregated Ethernet Interfaces

To specify aggregated Ethernet interfaces, include the **vlan-tagging** statement at the **[edit interfaces aex]** hierarchy level:

```
[edit interfaces aex]
vlan-tagging;
```

You must also include the **vlan-id** statement:

```
vlan-id number;
```

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*]**

For more information about the **vlan-tagging** and **vlan-id** statements, see “[802.1Q VLANs Overview](#)” on page 19.

- Related Documentation**
- [vlan-id](#)
 - [vlan-tagging on page 199](#)

Configuring a Layer 3 Subinterface (CLI Procedure)

EX Series switches use Layer 3 subinterfaces to divide a physical interface into multiple logical interfaces, each corresponding to a VLAN. The switch uses the Layer 3 subinterfaces to route traffic between subnets.

To configure Layer 3 subinterfaces, you enable VLAN tagging and partition one or more physical ports into multiple logical interfaces, each corresponding to a VLAN ID.

Before you begin, make sure you set up your VLANs. See [Configuring VLANs for EX Series Switches \(CLI Procedure\)](#) or [Configuring VLANs for EX Series Switches \(J-Web Procedure\)](#).

To configure Layer 3 subinterfaces:

1. Enable VLAN tagging:

```
[edit interfaces interface-name]
user@switch# set vlan-tagging
```

2. Bind each VLAN ID to a logical interface:

```
[edit interfaces interface-name]  
user@switch# set unit logical-unit-number vlan-id vlan-id-number
```

**Related
Documentation**

- [Example: Configuring Layer 3 Subinterfaces for a Distribution Switch and an Access Switch on page 32](#)
- [Verifying That Layer 3 Subinterfaces Are Working on page 116](#)
- [Understanding Layer 3 Subinterfaces on page 12](#)

Configuring 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. Enabling unicast RPF on the switch interfaces filters traffic with source addresses that do not use the incoming interface as the best return path back to the source. When a packet comes into an interface, if that interface is not the best return path to the source, the switch discards the packet. If the incoming interface is the best return path to the source, the switch forwards the packet.



NOTE: On EX3200 and EX4200 switches, you can only enable unicast RPF globally, on all switch interfaces. You cannot enable unicast RPF on a per-interface basis.

Before you begin:

- On an EX8200 switch, ensure that the selected switch interface is symmetrically routed before you enable unicast RPF. A symmetrically routed interface is an interface that uses the same route in both directions between the source and the destination. Do not enable unicast RPF on asymmetrically routed interfaces. An asymmetrically routed interface uses different paths to send and receive packets between the source and the destination.
- On an EX3200 or EX4200 switch, ensure that *all* switch interfaces are symmetrically routed before you enable unicast RPF on an interface. When you enable unicast RPF on any interface, it is enabled globally on all switch interfaces. Do not enable unicast RPF on asymmetrically routed interfaces. An asymmetrically routed interface uses different paths to send and receive packets between the source and the destination.

To enable unicast RPF, configure it explicitly on a selected customer-edge interface:

[edit interfaces]

user@switch# **set ge-1/0/10 unit 0 family inet rpf-check**



BEST PRACTICE: On EX3200 and EX4200 switches, unicast RPF is enabled globally on *all* switch interfaces, regardless of whether you configure it explicitly on only one interface or only on some interfaces.

On EX3200 and EX4200 switches, we recommend that you enable unicast RPF explicitly on either all interfaces or only one interface. To avoid possible confusion, do not enable it on only some interfaces:

- Enabling unicast RPF explicitly on only one interface makes it easier if you choose to disable it in the future because you must explicitly disable unicast RPF on every interface on which you explicitly enabled it. If you explicitly enable unicast RPF on two interfaces and you disable it on only one interface, unicast RPF is still implicitly enabled globally on the switch. The drawback to this approach is that the switch displays the flag that indicates that unicast RPF is enabled only on interfaces on which unicast RPF is 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 to 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 39](#)
- [Verifying Unicast RPF Status on page 117](#)
- [Disabling Unicast RPF \(CLI Procedure\) on page 107](#)
- [Troubleshooting Unicast RPF on page 125](#)
- [Understanding Unicast RPF for EX Series Switches on page 13](#)

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 or EX4200 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 message **warning: statement not found** displays. 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 or EX4200 switch.

On EX8200 switches, the switch does not apply unicast RPF to an interface unless you explicitly enable that interface for unicast RPF.

To disable unicast RPF, delete its configuration from the interface:

[edit interfaces]

user@switch# **delete** ge-1/0/10 unit 0 family inet **rpf-check**



NOTE: On EX3200 and EX4200 switches, if you do not disable unicast RPF on every interface on which you explicitly enabled it, unicast RPF remains implicitly enabled on all interfaces.

**Related
Documentation**

- [Example: Configuring Unicast RPF on an EX Series Switch on page 39](#)
- [Verifying Unicast RPF Status on page 117](#)
- [Configuring Unicast RPF \(CLI Procedure\) on page 106](#)
- [Understanding Unicast RPF for EX Series Switches on page 13](#)

Configuring IP Directed Broadcast (CLI Procedure)

You can use IP directed broadcast on an EX Series switch to facilitate remote network management by sending broadcast packets to hosts on a specified subnet without broadcasting to the entire network. IP directed broadcast packets are broadcast on only the target subnet. The rest of the network treats IP directed broadcast packets as unicast packets and forwards them accordingly.

Before you begin to configure IP directed broadcast:

- Ensure that the subnet on which you want broadcast packets using IP direct broadcast is not directly connected to the Internet.
- Configure a routed VLAN interface (RVI) for the subnet that will be enabled for IP direct broadcast. See [Configuring Routed VLAN Interfaces \(CLI Procedure\)](#) or [Configuring VLANs for EX Series Switches \(J-Web Procedure\)](#).



NOTE: We recommend that you do not enable IP directed broadcast on subnets that have a direct connection to the Internet because of increased exposure to denial-of-service (DoS) attacks.

To enable IP directed broadcast for a specified subnet:

1. Add the target subnet's logical interfaces to the VLAN:

```
[edit interfaces]
user@switch# set ge-0/0/0.0 family ethernet-switching vlan members vl
user@switch# set ge-0/0/1.0 family ethernet-switching vlan members vl
```

2. Configure the Layer 3 interface on the VLAN that is the target of the IP directed broadcast packets:

```
[edit interfaces]
user@switch# set vlan.1 family inet address 10.1.2.1/24
```

3. Associate a Layer 3 interface with the VLAN:

```
[edit vlans]
user@switch# set vl l3-interface vlan.1
```

4. Enable the Layer 3 interface for the VLAN to receive IP directed broadcasts:

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

**Related
Documentation**

- [Example: Configuring IP Directed Broadcast on an EX Series Switch on page 43](#)
- [Understanding IP Directed Broadcast for EX Series Switches on page 17](#)

Tracing Operations of an Individual Router or Switch Interface

To trace the operations of individual router or switch interfaces, include the **traceoptions** statement at the **[edit interfaces *interface-name*]** hierarchy level:

```
[edit interfaces interface-name]
traceoptions {
  flag flag;
}
```

You can specify the following interface tracing flags:

- **all**—Trace all interface operations.
- **event**—Trace all interface events.
- **ipc**—Trace all interface interprocess communication (IPC) messages.
- **media**—Trace all interface media changes.

The interfaces **traceoptions** statement does not support a trace file. The logging is done by the kernel, so the tracing information is placed in the system **syslog** files.

**Related
Documentation**

- [Tracing Operations of the Interface Process on page 110](#)
- [Tracing Interface Operations Overview](#)

Tracing Operations of the Interface Process

To trace the operations of the router or switch interface process, dcd, include the **traceoptions** statement at the **[edit interfaces]** hierarchy level:

```
[edit interfaces]
traceoptions {
  file <filename> <files number> <match regular-expression> <size size> <world-readable |
    no-world-readable>;
  flag <flag> <disable>;
  no-remote-trace;
}
```

By default, interface process operations are placed in the file named dcd and three 1-MB files of tracing information are maintained.

You can specify the following flags in the **interfaces traceoptions** statement:

- **change-events**—Log changes that produce configuration events.
- **config-states**—Log the configuration state machine changes.
- **kernel**—Log configuration IPC messages to kernel.
- **kernel-detail**—Log details of configuration messages to kernel.

For general information about tracing, see the tracing and logging information in the [Junos OS System Basics Configuration Guide](#).

Related Documentation

- [Tracing Interface Operations Overview](#)
- [Tracing Operations of an Individual Router or Switch Interface on page 109](#)

Setting the Mode on an SFP+ Uplink Module (CLI Procedure)

SFP+ uplink modules are supported on EX3200 and EX4200 switches. You can use these uplink modules either for two SFP+ transceivers or four SFP transceivers. You configure the operating mode on the module to match the type of transceiver you want to use—that is, for SFP+ transceivers, you configure the 10-gigabit operating mode, and for SFP transceivers, you configure the 1-gigabit operating mode.

By default, the SFP+ uplink module operates in the 10-gigabit mode and supports only SFP+ transceivers. If you have not changed the module from the default setting and you want to use SFP+ transceivers, you do not need to configure the operating mode.

To set the operating mode of an SFP+ uplink module:

1. Change the operating mode to the appropriate mode for the transceiver type you want to use by using one of the following commands:

```
[edit]
user@switch# set chassis fpc 0 pic 1 sfpplus pic-mode 1g
```

```
[edit]
```

```
user@switch# set chassis fpc 0 pic 1 sfpplus pic-mode 10g
```

2. If the switch is running:

- Junos OS Release 10.1 or later, the changed operating mode takes effect immediately unless a port on the SFP+ uplink module is a Virtual Chassis port (VCP). If any port on the SFP+ uplink module is a VCP, the changed operating mode does not take effect until the next reboot of the switch.



NOTE: During the operating mode change, the Packet Forwarding Engine is restarted. In a Virtual Chassis configuration, this means that the Flexible PIC Concentrator connection with the master is dropped and then reconnected.

- Junos OS Release 10.0 or earlier, reboot the switch.

You can see whether the operating mode has been changed to the new mode you configured by issuing the **show chassis pic fpc-slot slot-number pic-slot 1** command.

Related Documentation

- Uplink Modules in EX3200 Switches
- Uplink Modules in EX4200 Switches
- Optical Interface Support in EX3200 Switches
- Optical Interface Support in EX4200 Switches

Configuring the Media Type on Dual-Purpose Uplink Ports (CLI Procedure)

EX2200-C switches provide two dual-purpose uplink ports. Each dual uplink port is a single interface that offers a choice of two connections: an RJ-45 connection for a copper Ethernet cable and an SFP connection for a fiber-optic Ethernet cable. You can choose to use either connection, but only one connection can be active at a time.

By default, if you plug a transceiver into the SFP connector, the port becomes a fiber-optic Gigabit Ethernet port, even if a copper Ethernet cable is plugged into the RJ-45 connection as well. If a transceiver is not plugged into the SFP connector, the port defaults to a copper 10/100/1000 Ethernet port.

You can constrain the use of the port to one connection type by configuring the media type for the port to be either copper or fiber. When you configure a media type on the port, the port will no longer accept the alternate connection type. For example, if you configure the uplink port as a fiber port and then plug a copper Ethernet cable into the RJ-45 connector, the interface will not come up.

To configure the media type for an uplink port:

```
user@switch# set interfaces interface-name media-type media-type
```

For example, to set the media type for uplink port **ge-0/1/0** to copper:

```
user@switch# set interfaces ge-0/1/0 media-type copper
```



.....

NOTE: When you change the media type setting for a dual-purpose uplink port, it can take up to 6 seconds for the interface to appear in operational commands.

.....

**Related
Documentation**

- [EX2200 Switches Hardware Overview](#)

CHAPTER 4

Verifying Interfaces

- [Monitoring Interface Status and Traffic on page 113](#)
- [Verifying the Status of a LAG Interface on page 114](#)
- [Verifying That LACP Is Configured Correctly and Bundle Members Are Exchanging LACP Protocol Packets on page 115](#)
- [Verifying That Layer 3 Subinterfaces Are Working on page 116](#)
- [Verifying Unicast RPF Status on page 117](#)
- [Verifying IP Directed Broadcast Status on page 119](#)

Monitoring Interface Status and Traffic

Purpose 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.

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 FPC** list. Details such as the admin status and link status are displayed in the table.



NOTE: By default, the details of the first member in the **Port for FPC** drop-down list 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. Also, click on 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).

- Related Documentation**
- [Configuring Gigabit Ethernet Interfaces \(J-Web Procedure\) on page 51](#)
 - [Configuring Gigabit Ethernet Interfaces \(CLI Procedure\) on page 48](#)

Verifying the Status of a LAG Interface

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

Action Enter the following command:

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

- Related Documentation**
- [Configuring Aggregated Ethernet Interfaces \(CLI Procedure\) on page 97](#)
 - [Configuring Aggregated Ethernet Interfaces \(J-Web Procedure\) on page 98](#)

- [Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch on page 21](#)

Verifying That LACP Is Configured Correctly and Bundle Members Are Exchanging LACP Protocol Packets

Verify that LACP has been set up correctly and that the bundle members are transmitting LACP protocol packets.

1. [Verifying the LACP Setup on page 115](#)
2. [Verifying That LACP Packets Are Being Exchanged on page 115](#)

Verifying the LACP Setup

Purpose Verify that the LACP has been set up correctly.

Action TO verify that LACP has been enabled as active on one end:

```
user@switch>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 This example shows that LACP has been configured with one side as active and the other as passive. When LACP is enabled, one side must be set as active in order 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
- Verifying the Status of a LAG Interface

Verifying That Layer 3 Subinterfaces Are Working

Purpose After configuring Layer 3 subinterfaces, verify they are set up properly and transmitting data.

Action 1. Use the **show interfaces** command to determine if you successfully created the subinterfaces and the links are up:

```
user@switch> show interfaces interface-name terse
```

Interface	Admin	Link	Proto	Local	Remote
ge-0/0/0	up	up			
ge-0/0/0.0	up	up	inet	1.1.1.1/24	
ge-0/0/0.1	up	up	inet	2.1.1.1/24	
ge-0/0/0.2	up	up	inet	3.1.1.1/24	
ge-0/0/0.3	up	up	inet	4.1.1.1/24	
ge-0/0/0.4	up	up	inet	5.1.1.1/24	
ge-0/0/0.32767	up	up			

2. Use the **ping** command from a device on one subnet to an address on another subnet to determine if packets were transmitted correctly on the subinterface VLANs:

```
user@switch> ping ip-address
```

```

PING 1.1.1.1 (1.1.1.1): 56 data bytes
64 bytes from 1.1.1.1: icmp_seq=0 ttl=64 time=0.157 ms
64 bytes from 1.1.1.1: icmp_seq=1 ttl=64 time=0.238 ms
64 bytes from 1.1.1.1: icmp_seq=2 ttl=64 time=0.255 ms
64 bytes from 1.1.1.1: icmp_seq=3 ttl=64 time=0.128 ms
--- 1.1.1.1 ping statistics ---
4 packets transmitted, 4 packets received, 0% packet loss

```

Meaning The output confirms that the subinterfaces are created and the links are up.

- Related Documentation**
- [Configuring a Layer 3 Subinterface \(CLI Procedure\) on page 105](#)
 - [Example: Configuring Layer 3 Subinterfaces for a Distribution Switch and an Access Switch on page 32](#)

Verifying Unicast RPF Status

Purpose Verify that unicast reverse-path forwarding (RPF) is enabled and is working on the interface.

Action Use one of the **show interfaces *interface-name*** commands with either the **extensive** or **detail** options to verify that unicast RPF is enabled and working on the switch. The example below displays output from the **show interfaces ge- extensive** 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                                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 and EX4200 switches, unicast RPF is implicitly enabled on *all* switch interfaces, including aggregated Ethernet interfaces (also referred to as link aggregation groups or LAGs) 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 264](#)
 - [Example: Configuring Unicast RPF on an EX Series Switch on page 39](#)
 - [Configuring Unicast RPF \(CLI Procedure\) on page 106](#)
 - [Disabling Unicast RPF \(CLI Procedure\) on page 107](#)
 - [Troubleshooting Unicast RPF on page 125](#)

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 the following example.

- Related Documentation**
- [Configuring IP Directed Broadcast \(CLI Procedure\) on page 108](#)
 - [Example: Configuring IP Directed Broadcast on an EX Series Switch on page 43](#)

CHAPTER 5

Troubleshooting Interfaces

- [Troubleshooting Network Interfaces on EX3200 Switches on page 121](#)
- [Troubleshooting Network Interfaces on EX4200 Switches on page 122](#)
- [Troubleshooting an Aggregated Ethernet Interface on page 123](#)
- [Troubleshooting Interface Configuration and Cable Faults on page 124](#)
- [Troubleshooting Unicast RPF on page 125](#)
- [Troubleshooting Virtual Chassis Port Connectivity on an EX4200 Switch on page 125](#)
- [Diagnosing a Faulty Twisted-Pair Cable \(CLI Procedure\) on page 126](#)

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 121](#)
- [The interface on the port in which an SFP or SFP+ transceiver is installed in an SFP+ uplink module is down on page 122](#)

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 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.

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.

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 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.

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+ Uplink Module \(CLI Procedure\)”](#) on page 110.

- Related Documentation**
- [Troubleshooting Uplink Module Installation or Replacement on EX3200 Switches](#)
 - [Monitoring Interface Status and Traffic on page 113](#)
 - [Configuring Gigabit Ethernet Interfaces \(CLI Procedure\) on page 48](#)
 - [Configuring Gigabit Ethernet Interfaces \(J-Web Procedure\) on page 51](#)
 - [Removing a Transceiver from an EX Series 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 in an SFP+ uplink module is down on page 122](#)

The interface on the port in which an SFP or SFP+ transceiver is installed in an SFP+ uplink module is down

Problem The interface on the port in which an SFP or SFP+ transceiver is installed in an SFP+ uplink module installed in an EX4200 switch is down.

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+ Uplink Module \(CLI Procedure\)”](#) on page 110.

Related Documentation

- [Troubleshooting Virtual Chassis Port Connectivity on an EX4200 Switch](#) on page 125
- [Monitoring Interface Status and Traffic](#) on page 113
- [Configuring Gigabit Ethernet Interfaces \(CLI Procedure\)](#) on page 48
- [Configuring Gigabit Ethernet Interfaces \(J-Web Procedure\)](#) on page 51
- [Removing a Transceiver from an EX Series Switch](#)
- [Uplink Modules in EX4200 Switches](#)
- [EX Series Switches Interfaces Overview](#) on page 3

Troubleshooting an Aggregated Ethernet Interface

Troubleshooting issues for aggregated Ethernet interfaces:

- [Show Interfaces Command Shows the LAG is Down](#) on page 123
- [Logical Interface Statistics Do Not Reflect All Traffic](#) on page 123

Show Interfaces Command Shows the LAG is Down

Problem 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 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.

- Related Documentation**
- [Verifying the Status of a LAG Interface 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 21](#)
 - [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 27](#)

Troubleshooting Interface Configuration and Cable Faults

Troubleshooting interface configuration and connectivity on the EX Series switch:

1. [Interface Configuration or Connectivity Is Not Working on page 124](#)

Interface Configuration or Connectivity Is Not Working

Problem You encounter errors when you attempt to configure an interface on the switch, or the interface is exhibiting connectivity problems.

Solution Use the port troubleshooter feature in the J-Web interface to identify and rectify port configuration and connectivity related problems.

To use the J-Web interface port troubleshooter:

1. Select the option **Troubleshoot** from the main menu.
2. Click **Troubleshoot Port**. The Port Troubleshooting wizard is displayed. Click **Next**.
3. Select the ports to troubleshoot.
4. Select the test cases to be executed on the selected port. Click **Next**.

When the selected test cases are executed, the final result and the recommended action is displayed.

If there is a cable fault, the port troubleshooter displays details and the recommended action. For example, the cable must be replaced.

If the port configuration needs to be modified, the port troubleshooter displays details and the recommended action.

- Related Documentation**
- [Monitoring Interface Status and Traffic on page 113](#)
 - [Configuring Gigabit Ethernet Interfaces \(J-Web Procedure\) on page 51](#)
 - [Configuring Gigabit Ethernet Interfaces \(CLI Procedure\) on page 48](#)
 - [Connecting and Configuring an EX Series Switch \(CLI Procedure\)](#)

- [Connecting and Configuring an EX Series Switch \(J-Web Procedure\)](#)

Troubleshooting Unicast RPF

Troubleshooting issues for unicast reverse-path forwarding (RPF) on EX Series switches include:

1. [Legitimate Packets Are Discarded on page 125](#)

Legitimate Packets Are Discarded

Problem The switch filters valid packets from legitimate sources, which results in the switch's discarding packets that should be forwarded.

Solution The interface or interfaces on which legitimate packets are discarded are asymmetrically routed interfaces. An asymmetrically routed interface uses different paths to send and receive packets between the source and the destination, so the interface that receives a packet is not the same interface the switch uses to reply to the packet's source.

Unicast RPF works properly only on symmetrically routed interfaces. A symmetrically routed interface is an interface that uses the same route in both directions between the source and the destination. Unicast RPF filters packets by checking the forwarding table for the best return path to the source of an incoming packet. If the best return path uses the same interface as the interface that received the packet, the switch forwards the packet. If the best return path uses a different interface than the interface that received the packet, the switch discards the packet.



NOTE: On EX3200 and EX4200 switches, unicast RPF works properly only if all switch interfaces—including aggregated Ethernet interfaces (also referred to as link aggregation groups or LAGs) and routed VLAN interfaces (RVIs)—are symmetrically routed, because unicast RPF is enabled globally on all switch interfaces.

- Related Documentation**
- [Verifying Unicast RPF Status on page 117](#)
 - [Understanding Unicast RPF for EX Series Switches on page 13](#)

Troubleshooting Virtual Chassis Port Connectivity on an EX4200 Switch

This topic provides troubleshooting information for specific problems related to uplink module ports on EX4200 switches.

1. [Virtual Chassis port \(VCP\) connection does not work on page 126](#)

Virtual Chassis port (VCP) connection does not work

Problem The Virtual Chassis port (VCP) connection configured in an EX4200 switch does not work.

A port of the uplink module is set as a VCP.

Cause The uplink module installed in the switch was replaced.

Solution Set a port in the uplink module as a VCP. See [Setting an Uplink Port on an EX3300 or EX4200 Switch as a Virtual Chassis Port \(CLI Procedure\)](#).

- Related Documentation**
- [Monitoring Interface Status and Traffic on page 113](#)
 - [Configuring Gigabit Ethernet Interfaces \(CLI Procedure\) on page 48](#)
 - [Configuring Gigabit Ethernet Interfaces \(J-Web Procedure\) on page 51](#)
 - [Installing an Uplink Module in an EX4200 Switch](#)
 - [Removing a Transceiver from an EX Series Switch](#)
 - [Uplink Modules in EX4200 Switches](#)
 - [Understanding EX4200 and EX4500 Virtual Chassis Hardware Configurations](#)

Diagnosing a Faulty Twisted-Pair Cable (CLI Procedure)

Problem A 10/100/1000BASE-T Ethernet interface has connectivity problems that you suspect might be caused by a faulty cable.

Solution Use the time domain reflectometry (TDR) test to determine whether a twisted-pair Ethernet cable is faulty.

The TDR test:

- Detects and reports faults for each twisted pair in an Ethernet cable. Faults detected include open circuits, short circuits, and impedance mismatches.
- Reports the distance to fault to within 1 meter.
- Detects and reports pair swaps, pair polarity reversals, and excessive pair skew.

The TDR test is supported on the following switches and interfaces:

- EX2200, EX3200, EX3300, and EX4200 switches—RJ-45 network interfaces. The TDR test is not supported on management interfaces and SFP interfaces.
- EX6200 and EX8200 switches—RJ-45 network interfaces on line cards.



NOTE: We recommend running the TDR test on an interface when there is no traffic on the interface.

To diagnose a cable problem by running the TDR test:

1. Run the `request diagnostics tdr` command.

```
user@switch> request diagnostics tdr start interface ge-0/0/10
```

```
Interface TDR detail:
```

```
Test status : Test successfully executed ge-0/0/10
```

2. View the results of the TDR test with the `show diagnostics tdr` command.

```
user@switch> show diagnostics tdr interface ge-0/0/10
```

```
Interface TDR detail:
```

```
Interface name : ge-0/0/10
```

```
Test status : Passed
```

```
Link status : Down
```

```
MDI pair : 1-2
```

```
  Cable status : Normal
```

```
  Distance fault : 0 Meters
```

```
  Polartiy swap : N/A
```

```
  Skew time : N/A
```

```
MDI pair : 3-6
```

```
  Cable status : Normal
```

```
  Distance fault : 0 Meters
```

```
  Polartiy swap : N/A
```

```
  Skew time : N/A
```

```
MDI pair : 4-5
```

```
  Cable status : Open
```

```
  Distance fault : 1 Meters
```

```
  Polartiy swap : N/A
```

```
  Skew time : N/A
```

```
MDI pair : 7-8
```

```
  Cable status : Normal
```

```
  Distance fault : 0 Meters
```

```
  Polartiy swap : N/A
```

```
  Skew time : N/A
```

```
Channel pair : 1
```

```
  Pair swap : N/A
```

```
Channel pair : 2
```

```
  Pair swap : N/A
```

```
Downshift : N/A
```

3. Examine the **Cable status** field for the four MDI pairs to determine if the cable has a fault. In the preceding example, the twisted pair on pins 4 and 5 is broken or cut at approximately one meter from the **ge-0/0/10** port connection.



NOTE: The **Test Status** field indicates the status of the TDR test, not the cable. The value **Passed** means the test completed—it does not mean that the cable has no faults.

The following is additional information about the TDR test:

- The TDR test can take some seconds to complete. If the test is still running when you execute the `show diagnostics tdr` command, the **Test status** field displays **Started**. For example:

```
user@switch> show diagnostics tdr interface ge-0/0/22
```

```
Interface TDR detail:
```

```
Interface name      : ge-0/0/22
Test status         : Started
```

- You can terminate a running TDR test before it completes by using the **request diagnostics tdr abort interface *interface-name*** command. The test terminates with no results, and the results from any previous test are cleared.
- You can display summary information about the last TDR test results for all interfaces on the switch that support the TDR test by not specifying an interface name with the **show diagnostics tdr** command. For example:

```
user@switch> show diagnostics tdr
```

Interface	Test status	Link status	Cable status	Max distance fault
ge-0/0/0	Passed	UP	OK	0
ge-0/0/1	Not Started	N/A	N/A	N/A
ge-0/0/2	Passed	UP	OK	0
ge-0/0/3	Not Started	N/A	N/A	N/A
ge-0/0/4	Passed	UP	OK	0
ge-0/0/5	Passed	UP	OK	0
ge-0/0/6	Passed	UP	OK	0
ge-0/0/7	Not Started	N/A	N/A	N/A
ge-0/0/8	Passed	Down	OK	0
ge-0/0/9	Not Started	N/A	N/A	N/A
ge-0/0/10	Passed	Down	Fault	1
ge-0/0/11	Passed	UP	OK	0
ge-0/0/12	Not Started	N/A	N/A	N/A
ge-0/0/13	Not Started	N/A	N/A	N/A
ge-0/0/14	Not Started	N/A	N/A	N/A
ge-0/0/15	Not Started	N/A	N/A	N/A
ge-0/0/16	Not Started	N/A	N/A	N/A
ge-0/0/17	Not Started	N/A	N/A	N/A
ge-0/0/18	Not Started	N/A	N/A	N/A
ge-0/0/19	Passed	Down	OK	0
ge-0/0/20	Not Started	N/A	N/A	N/A
ge-0/0/21	Not Started	N/A	N/A	N/A
ge-0/0/22	Passed	UP	OK	0
ge-0/0/23	Not Started	N/A	N/A	N/A

**Related
Documentation**

- [Troubleshooting Interface Configuration and Cable Faults on page 124](#)
- [request diagnostics tdr on page 210](#)
- [show diagnostics tdr on page 212](#)

CHAPTER 6

Configuration Statements for Interfaces

- [\[edit chassis\] Configuration Statement Hierarchy on page 129](#)
- [\[edit interfaces\] Configuration Statement Hierarchy on page 130](#)

[\[edit chassis\] Configuration Statement Hierarchy](#)

```
chassis {
  aggregated-devices {
    ethernet {
      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 (slot-number | [list-of-slot-numbers]);
      member member-id {
        fpcs (slot-number | [list-of-slot-numbers]);
      }
    }
  }
  psu {
    redundancy {
      n-plus-n;
    }
  }
  redundancy {
```

```
        graceful-switchover;  
    }  
}
```

**Related
Documentation**

- [Configuring Aggregated Ethernet Interfaces \(CLI Procedure\) on page 97](#)
- [Upgrading Software Using Automatic Software Download on EX Series Switches](#)
- [Configuring the LCD Panel on EX Series Switches \(CLI Procedure\)](#)
- [Configuring Graceful Routing Engine Switchover in an EX3300, EX4200, or EX4500 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

```
interfaces {  
  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;
}
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);
    link-mode mode;
    (loopback | no-loopback);
    speed (auto-negotiation | speed);
  }
  (gratuitous-arp-reply | no-gratuitous-arp-reply);
  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;
}
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);
  link-mode mode;
  (loopback | no-loopback);
}

```

```
    speed (auto-negotiation | speed);
  }
  (gratuitous-arp-reply | no-gratuitous-arp-reply);
  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-id-number;
  }
  vlan-tagging;
}
lo0 {
  accounting-profile name;
  description text;
  disable;
  traceoptions {
    flag flag;
  }
  (traps | no-traps);
  unit logical-unit-number {
    accounting-profile name;
    bandwidth rate;
    description text;
    disable;
    family family-name {...}
    (traps | no-traps);
  }
}
me0 {
  accounting-profile name;
  description text;
  disable;
  (gratuitous-arp-reply | no-gratuitous-arp-reply);
  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;
}
vlan {
    accounting-profile name;
    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);
    }
}
traceoptions {
    file <filename> <files number> <match regular-expression> <size size>
    <world-readable | no-world-readable>;
    flag flag <disable>;
    no-remote-trace;
}
vme {
    accounting-profile name;
    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 {...}
        (traps | no-traps);
        vlan-id vlan-id-number;
    }
    vlan-tagging;
}
xe-fpc/pic/port {
    accounting-profile name;

```

```
description text;
disable;
ether-options {
  802.3ad {
    aex;
    (backup | primary);
    lacp {
      force-up;
    }
  }
  (flow-control | no-flow-control);
  link-mode mode;
  (loopback | no-loopback);
}
(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;
}
```

**Related
Documentation**

- [Configuring Gigabit Ethernet Interfaces \(CLI Procedure\) on page 48](#)
- [Configuring Aggregated Ethernet Interfaces \(CLI Procedure\) on page 97](#)
- [Configuring a Layer 3 Subinterface \(CLI Procedure\) on page 105](#)
- [Configuring Routed VLAN Interfaces \(CLI Procedure\)](#)
- [Configuring the Virtual Management Ethernet Interface for Global Management of an EX3300, EX4200, or EX4500 Virtual Chassis \(CLI Procedure\)](#)
- [EX Series Switches Interfaces Overview on page 3](#)
- [Junos OS Interfaces Fundamentals Configuration Guide](#)
- [Junos OS Ethernet Interfaces Configuration Guide](#)

802.3ad

Syntax	<pre>802.3ad { aex; (backup primary); lacp { force-up; } }</pre>
Hierarchy Level	[edit interfaces interface-name ether-options], [edit interfaces interface-range <i>name</i> ether-options]
Release Information	Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Configure membership in a link aggregation group (LAG).
Options	<ul style="list-style-type: none"> • aex—Name of the LAG. • backup—Designate the interface as the backup interface for link-protection mode. • primary—Designate the interface as the primary interface for link-protection mode. <p>The remaining statements are described separately.</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 21 • 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 27 • Configuring Aggregated Ethernet Interfaces (CLI Procedure) on page 97 • Configuring Aggregated Ethernet LACP (CLI Procedure) on page 101 • Understanding Aggregated Ethernet Interfaces and LACP on page 8 • Junos OS Ethernet Interfaces Configuration Guide

accounting-profile

Syntax	<code>accounting-profile <i>name</i>;</code>
Hierarchy Level	<code>[edit interfaces <i>interface-name</i>],</code> <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>],</code> <code>[edit interfaces interface-range <i>name</i>]</code>
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	<code>interface</code> —To view this statement in the configuration. <code>interface-control</code> —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 90• Applying an Accounting Profile to the Logical Interface on page 91

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 dlcidlcid-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 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.1 for QFX Series switches.

Description Configure the interface address.



NOTE: The vrrp High Availability functionality is not available for the QFX Series switches

Options *address*—Address of the interface.

The remaining statements are explained separately.



NOTE: The edit logical-systems hierarchy is not available on QFabric switches.

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
 - negotiate-address
 - unnumbered-address (Ethernet)
 - [Junos OS System Basics Configuration Guide](#)

aggregated-devices

Syntax	<pre> aggregated-devices { ethernet { device-count <i>number</i>; } } </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 21 • Configuring Aggregated Ethernet Interfaces (CLI Procedure) on page 97 • Understanding Aggregated Ethernet Interfaces and LACP on page 8 • Junos OS Ethernet Interfaces Configuration Guide

aggregated-ether-options

Syntax 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*;
 }

Hierarchy Level [edit [interfaces](#) aex]

Release Information Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Configure the aggregated Ethernet properties of a specific aggregated Ethernet interface.


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

- [Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch on page 21](#)
- [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 27](#)
- [Configuring Aggregated Ethernet Interfaces \(CLI Procedure\) on page 97](#)
- [Configuring Aggregated Ethernet LACP \(CLI Procedure\) on page 101](#)
- [Understanding Aggregated Ethernet Interfaces and LACP on page 8](#)
- [Junos OS Ethernet Interfaces Configuration Guide](#)

arp

Syntax	<code>arp <i>ip-address</i> (mac multicast-mac) <i>mac-address</i> publish;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i>], [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>]
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	For Ethernet, Fast Ethernet, and Gigabit Ethernet interfaces only, configure Address Resolution Protocol (ARP) table entries, mapping IP addresses to MAC addresses.
Options	<p><i>ip-address</i>—IP address to map to the MAC address. The IP address specified must be part of the subnet defined in the enclosing address statement.</p> <p>mac <i>mac-address</i>—MAC address to map to the IP address. Specify the MAC address as six hexadecimal bytes in one of the following formats: <i>nnnn.nnnn.nnnn</i> or <i>nn:nn:nn:nn:nn:nn</i>. For example, 0011.2233.4455 or 00:11:22:33:44:55.</p> <p>multicast-mac <i>mac-address</i>—Multicast MAC address to map to the IP address. Specify the multicast MAC address as six hexadecimal bytes in one of the following formats: <i>nnnn.nnnn.nnnn</i> or <i>nn:nn:nn:nn:nn:nn</i>. For example, 0011.2233.4455 or 00:11:22:33:44:55.</p> <p>publish—(Optional) Have the router or switch reply to ARP requests for the specified IP address. If you omit this option, the router or switch uses the entry to reach the destination but does not reply to ARP requests.</p>
<div>  <p>NOTE: The edit <code>logical-systems</code> hierarchy is not available on QFabric switches.</p> </div>	
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 Static ARP Table Entries on page 94 • Configuring Static ARP Entries

auto-negotiation

Syntax	(auto-negotiation no-auto-negotiation) <remote-fault (local-interface-online local-interface-offline) >;
Hierarchy Level	[edit interfaces <i>interface-name</i> ether-options], [edit interfaces <i>interface-name</i> gigether-options], [edit interfaces <i>ge-pim</i> /0/0 switch-options switch-port <i>port-number</i>]
Release Information	Statement introduced in Junos OS Release 7.6. Statement introduced in Junos OS Release 8.4 for J Series Services Routers. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	<p>For Gigabit Ethernet interfaces on M Series, MX Series, T Series, and TX Matrix routers, explicitly enable autonegotiation and remote fault. For EX Series switches and J Series Services Routers, explicitly enable autonegotiation only.</p> <ul style="list-style-type: none"> • auto-negotiation—Enables autonegotiation. This is the default. • no-auto-negotiation—Disable autonegotiation. When autonegotiation is disabled, you must explicitly configure the link mode and speed. <p>When you configure Tri-Rate Ethernet copper interfaces to operate at 1 Gbps, autonegotiation must be enabled.</p> <div style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;"> <p> NOTE: On EX Series switches, an interface configuration that disables autonegotiation and manually sets the link speed to 1 Gbps is accepted when you commit the configuration; however, if the interface you are configuring is a Tri-Rate Ethernet copper interface, the configuration is ignored as invalid and autonegotiation is enabled by default.</p> <p>To correct the invalid configuration and disable autonegotiation:</p> <ol style="list-style-type: none"> 1. Delete the no-auto-negotiation statement and commit the configuration. 2. Set the link speed to 10 or 100 Mbps, set no-auto-negotiation, and commit the configuration. </div> <p>On J Series Services Routers with universal Physical Interface Modules (uPIMs) and on EX Series switches, if the link speed and duplex mode are also configured, the interfaces use the values configured as the desired values in the negotiation. If autonegotiation is disabled, the link speed and link mode must be configured.</p>
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, and TX Matrix 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 J Series Services Router Switching Interfaces
- [Configuring Gigabit Ethernet Interfaces \(CLI Procedure\) on page 48](#)
- [Configuring Gigabit Ethernet Interfaces \(CLI Procedure\) on page 48](#)

bandwidth

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 68

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, nor can you specify a broadcast address.
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>  <p>NOTE: The edit logical-systems hierarchy is not available on QFabric switches.</p> </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 66

chassis

```
Syntax  chassis {
          aggregated-devices {
            ethernet {
              device-count number;
            }
          }
          auto-image-upgrade;
          fpc slot {
            pic pic-number {
              sfplusplus {
                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 (slot-number | [list-of-slot-numbers]);
              member member-id {
                fpcs (slot-number | [list-of-slot-numbers]);
              }
            }
          }
          psu {
            redundancy {
              n-plus-n;
            }
          }
          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 Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

- Related Documentation**
- [Configuring Aggregated Ethernet Interfaces \(CLI Procedure\) on page 97](#)
 - Upgrading Software Using Automatic Software Download on EX Series Switches
 - Configuring the LCD Panel on EX Series Switches (CLI Procedure)
 - Configuring Graceful Routing Engine Switchover in an EX3300, EX4200, or EX4500 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	<code>description text;</code>
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 on page 62 • Adding a Logical Unit Description to the Configuration on page 63 • Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 48 • Enabling and Disabling Insertion of Option 82 Information

device-count

Syntax	<code>device-count <i>number</i>;</code>
Hierarchy Level	[edit chassis aggregated-devices ethernet]
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 21• Configuring Aggregated Ethernet Interfaces (CLI Procedure) on page 97• Junos OS Network Interfaces Configuration Guide

disable (Interface)

Syntax	disable;
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	Disable a physical or a logical interface, effectively unconfiguring it.



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.



NOTE: 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 (GE) PICs with SFP and XFP transceivers do support it and the laser will be turned off when the interface is disabled.



NOTE: When you disable or deactivate an interface, then all the references made to the deactivated interface must be removed from the routing instance.



WARNING: Do not stare into the laser beam or view it directly with optical instruments even if the interface has been 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"> • Disabling a Physical Interface on page 64 • Disabling a Logical Interface on page 65
------------------------------	---

ether-options

Syntax Gigabit Ethernet interfaces:

```
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 (speed | auto-negotiation);  
}
```

10-Gigabit Ethernet interfaces:

```
ether-options {  
  802.3ad {  
    aex;  
    (backup | primary);  
    lacp {  
      force-up;  
    }  
  }  
  (flow-control | no-flow-control);  
  (loopback | no-loopback);  
}
```

Hierarchy Level [edit [interfaces interface-name](#)],
[edit interfaces interface-range *name*]

Release Information Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Configure Ethernet properties for a Gigabit Ethernet interface or a 10-Gigabit Ethernet interface on an EX Series switch.

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 48](#)
- [Configuring Gigabit Ethernet Interfaces \(J-Web Procedure\) on page 51](#)
- [Understanding Aggregated Ethernet Interfaces and LACP on page 8](#)
- [EX Series Switches Interfaces Overview on page 3](#)
- [Junos OS Ethernet Interfaces Configuration Guide](#)

ethernet

Syntax	ethernet { <code>device-count</code> <i>number</i> ; }
Hierarchy Level	[edit <code>chassis aggregated-devices</code>]
Release Information	Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Configure properties for Ethernet aggregated devices on the switch. The remaining statement is 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"> • Configuring Aggregated Ethernet Interfaces (CLI Procedure) on page 97 • Junos OS Ethernet 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 11.1 for QFX Series switches.
Description	For interfaces that carry IP version 6 (IPv6) traffic, automatically generate the host number portion of interface addresses. Not supported on QFX Series switches.
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 66

family (for EX Series switches)

Syntax	family ccc on page 152 family ethernet-switching on page 152 family inet on page 152 family inet6 on page 152 family iso on page 153 family mpls on page 153
family ccc	family ccc;
family ethernet-switching	family ethernet-switching { filter { input <i>filter-name</i> ; output <i>filter-name</i> ; } native-vlan-id <i>vlan-id</i> ; port-mode <i>mode</i> ; reflective-relay; vlan { members [(all <i>names</i> <i>vlan-ids</i>)]; } }
family inet	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> ; } } filter { input <i>filter-name</i> ; output <i>filter-name</i> ; } mtu <i>bytes</i> ; no-redirects ; no-neighbor-learn; primary; rpf-check ; targeted-broadcast ; }
family inet6	family inet6 { address <i>address</i> {

```

eui-64;
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;
}

```

```

family iso family iso {
    address interface-address;
    mtu bytes;
}

```

```

family mpls family mpls {
    mtu bytes;
}

```

Hierarchy Level [edit **interfaces** *interface-name* **unit** *logical-unit-number*],
[edit **interfaces** *interface-range name* **unit** *logical-unit-number*]

Release Information Statement introduced in Junos OS Release 9.0 for EX Series switches, including options **ethernet-switching**, **inet**, and **iso**.

Option **inet6** introduced in Junos OS Release 9.3 for EX Series switches.

Options **ccc** and **mpls** introduced in Junos OS Release 9.5 for EX Series switches.

Description Configure protocol family information for the logical interface on the switch.

Default Interfaces on EX2200, EX3200, EX3300, EX4200, and EX4500 switches are set to **family ethernet-switching** by the default factory configuration. If you are going to change the family setting for an interface, you might have to delete this default setting or any user-configured family setting before you change the setting to another family type.

EX6200 and EX8200 switch interfaces do not have a default family setting.

You must configure a logical interface to be able to use the physical device.

Options See [Table 26 on page 154](#) for protocol families available on the switch interfaces. Different protocol families support different subsets of the interfaces types on the switch.

Interface types on the switch are:

- Aggregated Ethernet (**ae**)
- Gigabit Ethernet (**ge**)
- Interface-range configuration (**interface-range**)
- Loopback (**lo0**)
- Management Ethernet (**me0**)
- Routed VLAN interface (RVI) (**vlan**)
- Virtual management Ethernet (**vme**)
- 10-Gigabit Ethernet (**xe**)

If you are using an interface range, the supported protocol families are the ones supported by the interface types that compose the range.

Not all interface types support all **family** substatements. Check your switch CLI for supported substatements for a particular protocol family configuration.

Table 26: Protocol Families and Supported Interface Types

Family	Description	Supported Interface Types						
		ae	ge	lo0	me0	vlan	vme	xe
ccc	Circuit cross-connect protocol family	✓*	✓					✓
ethernet-switching	Ethernet switching protocol family	✓	✓		✓			✓
inet	IPv4 protocol family	✓	✓	✓	✓	✓	✓	✓
inet6	IPv6 protocol family	✓	✓	✓	✓	✓	✓	✓
iso	Junos OS protocol family for IS-IS traffic	✓	✓	✓	✓	✓	✓	✓
mpls	MPLS protocol family	✓	✓	✓	✓		✓	✓

*Supported on EX8200 switches only

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**

- [Example: Configuring MPLS on EX Series Switches](#)
- [Configuring Gigabit Ethernet Interfaces \(CLI Procedure\) on page 48](#)
- [Configuring Aggregated Ethernet Interfaces \(CLI Procedure\) on page 97](#)
- [Configuring Routed VLAN Interfaces \(CLI Procedure\)](#)
- [*Junos OS Ethernet Interfaces Configuration Guide*](#)
- [*Junos OS Interfaces Fundamentals Configuration Guide*](#)

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	[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>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	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.
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">• Applying a Filter to an Interface• Configuring Firewall Filters (CLI Procedure)• Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 48• Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches• Junos OS Services Interfaces Configuration Guide• Junos OS Routing Policy Configuration Guide• Junos OS System Basics Configuration Guide

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> gigether-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.
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.
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 65 • Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 48

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 48 • Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 51 • 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>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 in EX Series switches.
Description	For Ethernet interfaces, enable updating of the ARP cache for replies received in response to gratuitous ARP requests.
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">• Configuring Gratuitous ARP on page 93• no-gratuitous-arp-request

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);`
 `link-mode mode;`
 `(loopback | no-loopback);`
 `speed (auto-negotiation | speed);`
 `}`
 `(gratuitous-arp-reply | no-gratuitous-arp-reply);`
 `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-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 use 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 interface—To view this statement in the configuration.
Level interface-control—To add this statement to the configuration.

- Related Documentation**
- [Configuring Gigabit Ethernet Interfaces \(CLI Procedure\) on page 48](#)
 - [Understanding Interface Ranges on EX Series Switches on page 10](#)
 - [EX Series Switches Interfaces Overview on page 3](#)
 - [Junos OS Interfaces Fundamentals Configuration Guide](#)

interfaces (for EX Series switches)

Syntax [interfaces ae on page 161](#)
[interfaces ge on page 161](#)
[interfaces interface-range on page 163](#)
[interfaces lo0 on page 163](#)
[interfaces me0 on page 164](#)
[interfaces traceoptions on page 164](#)
[interfaces vlan on page 164](#)
[interfaces vme on page 165](#)
[interfaces xe on page 166](#)

```

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);
link-mode mode;
(loopback | no-loopback);
speed (auto-negotiation | speed);
}
(gratuitous-arp-reply | no-gratuitous-arp-reply);
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);
        link-mode mode;
        (loopback | no-loopback);
        speed (auto-negotiation | speed);
    }
    (gratuitous-arp-reply | no-gratuitous-arp-reply);
    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;
    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);
    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);
    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);  
    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;  
}
```

```
interfaces xe xe-fpc/pic/port {  
    accounting-profile name;  
    description text;  
    disable;  
    ether-options {  
        802.3ad {  
            aex;  
            (backup | primary);  
            lacp {  
                force-up;  
            }  
        }  
        (flow-control | no-flow-control);  
        link-mode mode;  
        (loopback | no-loopback);  
    }  
    (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;  
}
```

Hierarchy Level [\[edit\]](#)

Release Information Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Configure interfaces on EX Series switches.

Options See [Table 27 on page 167](#) for the interface types and protocol-family options supported on the switch. Different protocol families support different subsets of the interface types on the switch. See the [family](#) statement for syntax of the protocol families supported for switch interfaces.

Not all interface types support all **family** substatements. Check your switch CLI for supported substatements for a particular protocol family configuration.

Table 27: Interface Types and Their Supported Protocol Families

Interface Type	Description	Supported Protocol Families					
		ccc	ethernet-switching	inet	inet6	iso	mpls
ae	Aggregated Ethernet interface (also referred to as a link aggregation group [LAG])	✓*	✓	✓	✓	✓	✓
ge	Gigabit Ethernet interface	✓	✓	✓	✓	✓	✓
lo0	Loopback interface			✓	✓	✓	✓
me0	Management Ethernet interface		✓	✓	✓	✓	✓
vlan	Routed VLAN interface (RVI)			✓	✓	✓	
vme	Virtual management Ethernet interface			✓	✓	✓	✓
xe	10-Gigabit Ethernet interface	✓	✓	✓	✓	✓	✓
interface-range	Interface-range configuration	Supported protocol families are the ones supported by the interface types that compose the range.					
*Supported on EX8200 switches only							

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 48](#)
 - [Configuring Aggregated Ethernet Interfaces \(CLI Procedure\) on page 97](#)
 - [Configuring a Layer 3 Subinterface \(CLI Procedure\) on page 105](#)
 - [Configuring Routed VLAN Interfaces \(CLI Procedure\)](#)
 - [Configuring the Virtual Management Ethernet Interface for Global Management of an EX3300, EX4200, or EX4500 Virtual Chassis \(CLI Procedure\)](#)
 - [EX Series Switches Interfaces Overview on page 3](#)
 - [Junos OS Interfaces Fundamentals Configuration Guide](#)
 - [Junos OS Ethernet Interfaces Configuration Guide](#)



lACP (802.3ad)

Syntax	<code>lACP { force-up; }</code>
Hierarchy Level	[edit interfaces interface-name ether-options 802.3ad]
Release Information	Statement introduced in Junos OS Release 10.0 for EX Series switches.
Description	Configure the Link Aggregation Control Protocol (LACP) parameters for 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">• Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch on page 21• 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 27• Configuring Aggregated Ethernet Interfaces (CLI Procedure) on page 97• Configuring Aggregated Ethernet LACP (CLI Procedure) on page 101• Understanding Aggregated Ethernet Interfaces and LACP on page 8• Junos OS Ethernet Interfaces Configuration Guide

lACP (Aggregated Ethernet)

Syntax	<pre>lACP { (active passive); admin-key <i>key</i>; link-protection { disable; (revertive non-revertive); } periodic <i>interval</i>; system-id <i>mac-address</i>; system-priority <i>priority</i>; }</pre>
Hierarchy Level	[edit interfaces <i>aex</i> aggregated-ether-options]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	For aggregated Ethernet interfaces only, configure Link Aggregation Control Protocol (LACP).
Default	If you do not specify LACP as either active or passive, LACP remains passive.
Options	<ul style="list-style-type: none"> • active—Initiate transmission of LACP packets. • passive—Respond to LACP packets. <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"> • Configuring Aggregated Ethernet LACP • Configuring Aggregated Ethernet LACP (CLI Procedure) on page 101 • 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 27

link-mode

Syntax	link-mode <i>mode</i> (automatic full-duplex half-duplex);
Hierarchy Level	[edit interfaces <i>interface-name</i>], [edit interfaces <i>interface-name</i> ether-options], [edit interfaces ge-pim/0/0 switch-options switch-port <i>port-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	Set the device's link connection characteristic.
Options	<p><i>mode</i>—Link characteristics:</p> <ul style="list-style-type: none"> • automatic—Link mode is negotiated. This is the default for EX Series switches. • full-duplex—Connection is full duplex. • half-duplex—Connection is half duplex. <p>Default: Fast Ethernet interfaces, except the J Series ePIM Fast Ethernet interfaces, can operate in either full-duplex or half-duplex mode. The router's management Ethernet interface, fxp0 or em0, the built-in Fast Ethernet interfaces on the FIC (M7i router), and the Gigabit Ethernet ports on J Series Services Routers with uPIMs installed and configured for access switching mode autonegotiate whether to operate in full-duplex or half-duplex mode. Unless otherwise noted here, all other interfaces operate only in full-duplex mode.</p>
	<p> NOTE: On J Series ePIM Fast Ethernet interfaces, if you specify half-duplex (or if full-duplex mode is not autonegotiated), the following message is written to the system log: "Half-duplex mode not supported on this PIC, forcing full-duplex mode."</p>
	<p> NOTE: On EX Series switches, if no-auto-negotiation is specified in [edit interfaces <i>interface-name</i> ether-options], you can select only full-duplex or half-duplex. If auto-negotiation is specified, you can select any mode.</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"> • Configuring the Link Characteristics on Ethernet Interfaces • Understanding Management Ethernet Interfaces

link-protection

Syntax	<pre>link-protection { disable; (revertive non-revertive); }</pre>
Hierarchy Level	[edit interfaces aex aggregated-ether-options] [edit interfaces aex aggregated-ether-options <i>lcp</i>]
Release Information	Statement introduced in Junos OS Release 8.3. Statement introduced in Junos OS Release 9.0 for EX Series switches. Support for disable , revertive , and non-revertive statements added in Junos OS Release 9.3.
Description	<p>On the router, for aggregated Ethernet interfaces only, configure link protection. In addition to enabling link protection, a primary and a secondary (backup) link must be configured to specify what links egress traffic should traverse. To configure primary and secondary links on the router, include the primary and backup statements at the [edit interfaces <i>ge-fpc/pic/port</i> gigether-options 802.3ad aex] hierarchy level or the [edit interfaces <i>fe-fpc/pic/port</i> fastether-options 802.3ad aex] hierarchy level.</p> <p>To configure those links on the switch, configure those statements at the [edit interfaces <i>ge-fpc/pic/port</i> ether-options 802.3ad aex] hierarchy level or at the [edit interfaces <i>xe-fpc/pic/port</i> ether-options 802.3ad aex] hierarchy level.</p>
Options	The 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"> • Configuring Aggregated Ethernet Link Protection on page 102

link-speed (Aggregated Ethernet)

Syntax	link-speed <i>speed</i> ;
Hierarchy Level	[edit interfaces aex aggregated-ether-options], [edit interfaces interface-range <i>name</i> aggregated-ether-options], [edit interfaces interface-range <i>name</i> aggregated-sonet-options]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	For aggregated Ethernet interfaces only, set the required link speed.
Options	<p>speed—For aggregated Ethernet links, you can specify speed 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).</p> <p>Aggregated Ethernet links on the M120 router can have one of the following speed values:</p> <ul style="list-style-type: none">• 100m—Links are 100 Mbps.• 10g—Links are 10 Gbps.• 1g—Links are 1 Gbps.• oc192—Links are OC192 or STM64c. <p>Aggregated Ethernet links on EX Series switches can be configured to operate at one of the following speed values:</p> <ul style="list-style-type: none">• 10m• 100m• 1g• 10g
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 Aggregated Ethernet Link Speed on page 103• Configuring Aggregated Ethernet Interfaces (CLI Procedure) on page 97• Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch on page 21

loopback (Aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet)

Syntax	(loopback no-loopback);
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> gigether-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 for EX Series switches.
Description	For aggregated Ethernet, Fast Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces, enable or disable loopback mode.



NOTE: By default, local aggregated Ethernet, Fast Ethernet, Tri-Rate Ethernet copper, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces connect to a remote system.

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 Ethernet Loopback Capability on page 92

media-type

Syntax	<code>media-type (copper fiber);</code>
Hierarchy Level	[edit interfaces <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 11.3 for EX Series switches.
Description	(EX2200-C switch only) Configure the media type for a dual-purpose uplink port (one RJ-45 port and one SFP port) on an EX2200 switch. If you use the media-type for a dual-purpose uplink port, the alternate media type cannot be used with the port.
Default	When media-type is not set, the port accepts either type of connection. The media type is fiber if a transceiver is installed in the SFP connection. If no transceiver is installed, the media type is copper.
Options	copper —The dual-purpose uplink port accepts only a 10/100/1000BASE-T copper connection. fiber —The dual-purpose uplink port accepts only an SFP fiber connection.
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 Media Type on Dual-Purpose Uplink Ports (CLI Procedure) on page 111

member

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 48• Understanding Interface Ranges on EX Series Switches on page 10• EX Series Switches Interfaces Overview on page 3• Junos OS Interfaces Fundamentals Configuration Guide

members

Syntax `members [(all | names | vlan-ids)];`

Hierarchy Level `[edit interfaces interface-name unit logical-unit-number family ethernet-switching vlan]`

Release Information Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement updated with enhanced ? (CLI completion feature) functionality in Junos OS Release 9.5 for EX Series switches.

Description For trunk interfaces, configure the VLANs that can carry traffic.



TIP: To display a list of all configured VLANs on the system, including VLANs that are configured but not committed, type ? after `vlan` or `vlan`s in your configuration mode command line. Note that only one VLAN is displayed for a VLAN range.



NOTE: The number of VLANs supported per switch varies for each model. Use the configuration-mode command `set vlans id vlan-id ?` to determine the maximum number of VLANs allowed on a switch. You cannot exceed this VLAN limit because each VLAN is assigned an ID number when it is created. You can, however, exceed the recommended VLAN member maximum. To determine the maximum number of VLAN members allowed on a switch, multiply the VLAN maximum for the switch times 8 ($\text{vmember limit} = \text{vlan max} * 8$).

If a switch configuration exceeds the recommended VLAN member maximum, you see a warning message when you commit the configuration. If you ignore the warning and commit such a configuration, the configuration succeeds but you run the risk of crashing the Ethernet switching process (`eswd`) due to memory allocation failure.

Options `all`—Specifies that this trunk interface is a member of all the VLANs that are configured on this switch. When a new VLAN is configured on the switch, this trunk interface automatically becomes a member of the VLAN.



NOTE: Since VLAN members are limited, specifying `all` could cause the number of VLAN members to exceed the limit at some point.

names—Name of one or more VLANs. VLAN IDs are applied automatically in this case.



NOTE: all cannot be a VLAN name.

vlan-ids—Numeric identifier of one or more VLANs. For a series of tagged VLANs, specify a range; for example, 10-20 or 10-20 23 27-30.



NOTE: Each configured VLAN must have a specified VLAN ID to successfully commit the configuration; otherwise, the configuration commit fails.

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"> • show ethernet-switching interfaces on page 217 • show vlans • Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch • Example: Connecting an Access Switch to a Distribution Switch • Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 48 • Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 51 • Configuring VLANs for EX Series Switches (CLI Procedure) • Creating a Series of Tagged VLANs (CLI Procedure) • Understanding Bridging and VLANs on EX Series Switches • Junos OS Ethernet Interfaces 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 48• Understanding Interface Ranges on EX Series Switches on page 10• EX Series Switches Interfaces Overview on page 3• Junos OS Interfaces Fundamentals Configuration Guide

minimum-links

Syntax	<code>minimum-links <i>number</i>;</code>
Hierarchy Level	[edit interfaces <i>aex</i> aggregated-ether-options], [edit interfaces <i>aex</i> aggregated-sonet-options], [edit interfaces <i>interface-name</i> mlfr-uni-nni-bundle-options], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit interfaces interface-range <i>range</i> aggregated-ether-options], [edit interfaces interface-range <i>range</i> aggregated-sonet-options], [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	For aggregated Ethernet, SONET/SDH, multilink, link services, and voice services interfaces only, set the minimum number of links that must be up for the bundle to be labeled up.
Options	<i>number</i> —Number of links. Range: 1 through 8 (1 through 16 for Ethernet and SONET interfaces on the MX Series, M320, M120, T Series, or TX Matrix routers, and 1 through 12 for EX8200 switches) Default: 1
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 Aggregated Ethernet Minimum Links on page 104• Configuring Aggregated SONET/SDH Minimum Links• Configuring Aggregated Ethernet Interfaces (CLI Procedure) on page 97• Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch on page 21• Junos OS Services Interfaces Configuration Guide

mtu

Syntax	<code>mtu bytes;</code>
Hierarchy Level	<code>[edit interfaces <i>interface-name</i>],</code> <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>],</code> <code>[edit interfaces interface-range <i>name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>]</code>
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>Specify the maximum transmission unit (MTU) size for the media or protocol. The default MTU size depends on the device type. Changing the media MTU or protocol MTU causes an interface to be deleted and added again.</p> <p>To route jumbo data packets on the RVI on EX Series switches, you must configure the jumbo MTU size on the member physical interfaces and also on the RVI itself (the VLAN interface).</p> <div style="margin-top: 20px;">  <p>CAUTION: For EX Series switches, setting or deleting the jumbo MTU size on the RVI (the vlan interface) while the switch is transmitting packets might result in dropped packets.</p> </div> <div style="margin-top: 20px;">  <p>NOTE: Not all devices allow you to set an MTU value, and some devices have restrictions on the range of allowable MTU values. You cannot configure an MTU for management Ethernet interfaces (fxp0, em0, or me0) or for loopback, multilink, and multicast tunnel devices.</p> </div> <p>For more information on configuring MTU for specific interfaces and router or switch combinations, see “Configuring the Media MTU” on page 69.</p>
Options	<p>bytes—MTU size.</p> <p>Range: 256 through 9192 bytes</p> <p>Default: 1500 bytes (INET, INET6, and ISO families), 1448 bytes (MPLS), 1514 bytes (EX Series interfaces)</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 the Media MTU on page 69 • Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 48

- [Configuring Routed VLAN Interfaces \(CLI Procedure\)](#)
- [Setting the Protocol MTU on page 80](#)

native-vlan-id

Syntax	<code>native-vlan-id <i>vlan-id</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit 0 family ethernet-switching]
Release Information	Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Configure the VLAN identifier to associate with untagged packets received on the interface.
Options	<i>vlan-id</i> —Numeric identifier of the VLAN. Range: 0 through 4095
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <code>show vlans</code>• show ethernet-switching interfaces on page 217• Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 48• Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 51• Understanding Bridging and VLANs on EX Series Switches• Junos OS Ethernet Interfaces Configuration Guide

no-redirects

Syntax	no-redirects;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</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>Do not send protocol redirect messages on the interface.</p> <p>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.</p>
Default	Interfaces send protocol redirect messages.
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">• Disabling the Transmission of Redirect Messages on an Interface on page 95• Junos OS System Basics Configuration Guide

periodic

Syntax	<code>periodic interval;</code>
Hierarchy Level	[edit interfaces aex aggregated-ether-options lACP], [edit interfaces interface-range <i>name</i> aggregated-ether-options lACP]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	For aggregated Ethernet interfaces only, configure the interval for periodic transmission of LACP packets.
Options	<i>interval</i> —Interval for periodic transmission of LACP packets. <ul style="list-style-type: none">• fast—Transmit packets every second.• slow—Transmit packets every 30 seconds. Default: fast
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 Aggregated Ethernet LACP• Configuring Aggregated Ethernet LACP (CLI Procedure) on page 101• Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch on page 21

pic

Syntax	<code>pic <i>pic-number</i> { sfpplus { pic-mode <i>mode</i>; } }</code>
Hierarchy Level	[edit <code>chassis fpc slot</code>]
Release Information	Statement introduced in Junos OS Release 9.4 for EX Series switches.
Description	Enable the specified port of the SFP+ uplink module to perform in the operating mode specified by <code>pic-mode</code> . The port is indicated by a Physical Interface Card (PIC) number.
Options	<p><code>pic-number</code>—Number of the PIC. For uplink ports in EX3200 and EX4200 switches, the PIC number is always 1.</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"> • Setting the Mode on an SFP+ Uplink Module (CLI Procedure) on page 110

pic-mode

Syntax	<code>pic-mode <i>mode</i>;</code>
Hierarchy Level	[edit <code>chassis fpc slot <i>pic</i> <i>pic-number</i> sfpplus</code>]
Release Information	Statement introduced in Junos OS Release 9.4 for EX Series switches.
Description	Configure the operating mode for the specified port on the SFP+ uplink module on an EX3200 or EX4200 switch.
Options	<p><code>mode</code>—Operating mode of the SFP+ uplink module:</p> <ul style="list-style-type: none"> • 1G—1-gigabit operating mode • 10G—10-gigabit operating mode
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"> • Setting the Mode on an SFP+ Uplink Module (CLI Procedure) on page 110

port-mode

Syntax	<code>port-mode mode;</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	Configure whether an interface on the switch operates in access, tagged-access, or trunk mode.
Default	All switch interfaces are in access mode.
Options	<p><i>mode</i>—Operating mode for an interface can be one of the following:</p> <ul style="list-style-type: none"> • access—In this mode, the interface can be in a single VLAN only. Access interfaces typically connect to single network devices such as PCs, printers, IP telephones, and IP cameras. • tagged-access—In this mode, the interface can accept tagged packets from one access device. Tagged-access interfaces typically connect to servers running Virtual machines using VEPA technology. • trunk—In this mode, the interface can be in multiple VLANs and accept tagged packets from multiple devices. Trunk interfaces typically connect to other switches and to routers on the LAN.




NOTE: The number of VLANs supported per switch varies for each model. Use the configuration-mode command `set vlans id vlan-id ?` to determine the maximum number of VLANs allowed on a switch. You cannot exceed this VLAN limit because each VLAN is assigned an ID number when it is created. You can, however, exceed the recommended VLAN member maximum. To determine the maximum number of VLAN members allowed on a switch, multiply the VLAN maximum for the switch times 8 ($\text{vmember limit} = \text{vlan max} * 8$).

If a switch configuration exceeds the recommended VLAN member maximum, you see a warning message when you commit the configuration. If you ignore the warning and commit such a configuration, the configuration succeeds but you run the risk of crashing the Ethernet switching process (`eswd`) due to memory allocation failure.


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: Connecting an Access Switch to a Distribution Switch • Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 48

- Configuring VLANs for EX Series Switches (CLI Procedure)
- [Junos OS Ethernet Interfaces Configuration Guide](#)

preferred

Syntax	preferred;
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	Configure this address to be the preferred address on the interface. If you configure more than one address on the same subnet, the preferred source address is chosen by default as the source address when you initiate frame transfers to destinations on the subnet.
<div>  <p>NOTE: The edit logical-systems hierarchy is not available on QFabric switches.</p> </div>	
Default	The lowest-numbered address on the subnet is the preferred address.
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 66

primary (Address on Interface)

Syntax	primary;
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	Configure this address to be the primary address of the protocol on the interface. If the logical unit has more than one address, the primary address is used by default as the source address when packet transfer originates from the interface and the destination address does not indicate the subnet.
<div> NOTE: The edit logical-systems hierarchy is not available on QFabric switches.</div>	
Default	For unicast traffic, the primary address is the lowest non-127 (in other words, non-loopback) preferred address on the unit.
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 66


proxy-arp

Syntax	<code>proxy-arp (restricted unrestricted);</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 9.6 for EX Series switches. restricted added in Junos OS Release 10.0 for EX Series switches.
Description	For Ethernet interfaces only, configure the router or switch to respond to any ARP request, as long as the router or switch has an active route to the ARP request's target address.
Default	Proxy ARP is not enabled. The router or switch responds to an ARP request only if the destination IP address is its own.
Options	<ul style="list-style-type: none"> • none—The router or switch responds to any ARP request for a local or remote address if the router or switch has a route to the target IP address. • restricted—(Optional) The router or switch responds to ARP requests in which the physical networks of the source and target are different and does not respond if the source and target IP addresses are in the same subnet. The router or switch must also have a route to the target IP address. • unrestricted—(Optional) The router or switch responds to any ARP request for a local or remote address if the router or switch has a route to the target IP address. <p>Default: unrestricted</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"> • Configuring Restricted and Unrestricted Proxy ARP on page 95 • Configuring Proxy ARP (CLI Procedure) • Example: Configuring Proxy ARP on an EX Series Switch • Configuring Gratuitous ARP on page 93

rpf-check

Syntax	rpf-check;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6]
Release Information	Statement introduced in Junos OS Release 9.3 for EX Series switches.
Description	<p>On EX3200 and EX4200 switches, enable a reverse-path forwarding (RPF) check on unicast traffic (except ECMP packets) on <i>all</i> ingress interfaces.</p> <p>On EX8200 switches, enable an RPF check on unicast traffic, including ECMP packets, on the selected ingress interface.</p>
Default	Unicast RPF is disabled on all 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">• Example: Configuring Unicast RPF on an EX Series Switch on page 39• Configuring Unicast RPF (CLI Procedure) on page 106• Disabling Unicast RPF (CLI Procedure) on page 107• Understanding Unicast RPF for EX Series Switches on page 13

sfpplus

Syntax	<code>sfpplus { pic-mode mode; }</code>
Hierarchy Level	[edit <code>chassis fpc slot pic pic-number</code>]
Release Information	Statement introduced in Junos OS Release 9.4 for EX Series switches.
Description	<p>Configure the operating mode for the specified port on the SFP+ uplink module on the EX3200 or EX4200 switch.</p> <p>The remaining statement is explained separately.</p>
Default	By default, the SFP+ uplink module operates in the 10-gigabit mode and supports SFP+ transceivers.
<div>  <p>NOTE: The SFP+ uplink module provides two ports for 10-gigabit small form-factor pluggable (SFP+) transceivers when configured to operate in 10-gigabit mode or four ports for 1-gigabit small form-factor pluggable (SFP) transceivers when configured to operate in 1-gigabit mode.</p> </div>	
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"> • Setting the Mode on an SFP+ Uplink Module (CLI Procedure) on page 110

speed

Syntax	<code>speed (auto-negotiation <i>speed</i>) ;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> ether-options]
Release Information	Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Configure the interface's speed.
Default	If the auto-negotiation statement at the [edit interfaces <i>interface-name</i> ether-options] hierarchy level is enabled, the auto-negotiation option is enabled by default.
Options	<ul style="list-style-type: none">• auto-negotiation—Automatically negotiate the speed based on the speed of the other end of the link. This option is available only when the auto-negotiation statement at the [edit interfaces <i>interface-name</i> ether-options] hierarchy level is enabled.• speed—Specify the interface speed. If the auto-negotiation statement at the [edit interfaces <i>interface-name</i> ether-options] hierarchy level is disabled, you must specify a specific value. This value sets the speed that is used on the link. If the auto-negotiation statement is enabled, you might want to configure a specific speed value to advertise the desired speed to the remote end.<ul style="list-style-type: none">• 10m—10 Mbps• 100m—100 Mbps• 1g—1 Gbps
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 48• Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 51• Junos OS Ethernet Interfaces Configuration Guide

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 an EX Series Switch on page 43• Configuring IP Directed Broadcast (CLI Procedure) on page 108• Understanding IP Directed Broadcast for EX Series Switches on page 17

traceoptions (Individual Interfaces)

Syntax	<pre>traceoptions { file <i>filename</i> <files <i>name</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i>; match; }</pre>
Hierarchy Level	[edit interfaces <i>interface-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	<p>Define tracing operations for individual interfaces.</p> <p>To specify more than one tracing operation, include multiple flag statements.</p> <p>The interfaces traceoptions statement does not support a trace file. The logging is done by the kernel, so the tracing information is placed in the system syslog file in the directory /var/log.</p>
Default	If you do not include this statement, no interface-specific tracing operations are performed.
Options	<p>file <i>name</i>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory /var/log. By default, interface process tracing output is placed in the file files <i>number</i>—(Optional) Maximum number of trace files. When a trace file named trace-file reaches its maximum size, it is renamed trace-file.0, then trace-file.1, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.dcd.</p> <p>match—(Optional) Regular expression for lines to be traced.</p> <p>no-world-readable—(Optional) Prevent any user from reading the log file.</p> <p>world-readable—(Optional) Allow any user to read the log file.</p> <p>size <i>size</i>—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named trace-file reaches this size, it is renamed trace-file.0. When the trace-file again reaches its maximum size, trace-file.0 is renamed trace-file.1 and trace-file is renamed trace-file.0. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.</p> <p>flag—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. The following are the interface-specific tracing options.</p> <ul style="list-style-type: none">• all—All interface tracing operations• event—Interface events• ipc—Interface interprocess communication (IPC) messages

- **media**—Interface media changes
- **q921**—Trace ISDN Q.921 frames
- **q931**—Trace ISDN Q.931 frames

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

Related Documentation • [Tracing Operations of an Individual Router or Switch Interface on page 109](#)

traceoptions (Interface Process)

Syntax	<pre>traceoptions { file <filename> <files number> <match regular-expression> <size size> <world-readable no-world-readable>; flag flag <disable>; no-remote-trace; }</pre>
Hierarchy Level	[edit interfaces]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Define tracing operations for the interface process (dcd).
Default	If you do not include this statement, no interface-specific tracing operations are performed.
Options	<p>disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as all.</p> <p>filename—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory /var/log. By default, interface process tracing output is placed in the file dcd.</p> <p>files number—(Optional) Maximum number of trace files. When a trace file named trace-file reaches its maximum size, it is renamed trace-file.0, then trace-file.1, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p>If you specify a maximum number of files, you also must specify a maximum file size with the size option.</p> <p>Range: 2 through 1000</p> <p>Default: 3 files</p> <p>flag—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags:</p> <ul style="list-style-type: none">• all• change-events—Log changes that produce configuration events• config-states—Log the configuration state machine changes• kernel—Log configuration IPC messages to kernel• kernel-detail—Log details of configuration messages to kernel <p>no-world-readable—(Optional) Disallow any user to read the log file.</p>

size size—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-file** again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

Syntax: **xk** to specify kilobytes, **xm** to specify megabytes, or **xg** to specify gigabytes

Range: 10 KB through the maximum file size supported on your router

Default: 1 MB

world-readable—(Optional) Allow any user to read the log file.

match regex—(Optional) Refine the output to include only those lines that match the given regular expression.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- [Tracing Operations of the Interface Process on page 110](#)

traps

Syntax	(traps no-traps);
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>], [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	Enable or disable the sending of Simple Network Management Protocol (SNMP) notifications when the state of the connection changes.
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"> • Enabling or Disabling SNMP Notifications on Physical Interfaces on page 97 • Enabling or Disabling SNMP Notifications on Logical Interfaces on page 96

unit

Syntax	<pre>unit <i>logical-unit-number</i> { accounting-profile <i>name</i>; bandwidth <i>rate</i>; description <i>text</i>; disable; family <i>family-name</i> {...} proxy-arp (restricted unrestricted); (traps no-traps); vlan-id <i>vlan-id-number</i>; }</pre>
Hierarchy Level	[edit interfaces <i>interface-name</i>], [edit interfaces interface-range <i>name</i>]
Release Information	Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Configure a logical interface on the physical device. You must configure a logical interface to be able to use the physical device.
Options	<p><i>logical-unit-number</i>—Number of the logical unit.</p> <p>Range: 0 through 16,384</p> <p>The remaining statements are explained separately.</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">• Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 48• Configuring Aggregated Ethernet Interfaces (CLI Procedure) on page 97• EX Series Switches Interfaces Overview on page 3• Junos OS Ethernet Interfaces Configuration Guide

vlan

Syntax	<pre>vlan { members [(all names vlan-ids)]; }</pre>
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"> • show ethernet-switching interfaces on page 217 • Example: Setting Up Bridging with Multiple VLANs for EX Series Switches • Configuring Routed VLAN Interfaces (CLI Procedure) • Understanding Bridging and VLANs on EX Series Switches • Junos OS Ethernet Interfaces Configuration Guide

vlan-id

Syntax	<code>vlan-id <i>vlan-id-number</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced in Junos OS Release 9.2 for EX Series switches.
Description	Bind an 802.1Q VLAN tag ID to a logical interface.



NOTE: The VLAN tag ID cannot be configured on logical interface unit 0. The logical unit number must be 1 or higher.

Options	<i>vlan-id-number</i> —A valid VLAN identifier. Range: 1 through 4094
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">• vlan-tagging on page 199• Example: Configuring Layer 3 Subinterfaces for a Distribution Switch and an Access Switch on page 32• Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 48• Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 51• Configuring a Layer 3 Subinterface (CLI Procedure) on page 105• Junos OS Ethernet Interfaces Configuration Guide

vlan-tagging

Syntax	vlan-tagging;
Hierarchy Level	[edit interfaces <i>interface-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	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	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">• 802.1Q VLANs Overview on page 19• vlan-id on page 198• Configuring a Layer 3 Subinterface (CLI Procedure) on page 105• Configuring Tagged Aggregated Ethernet Interfaces on page 105• Example: Configuring Layer 3 Subinterfaces for a Distribution Switch and an Access Switch on page 32

CHAPTER 7

Operational Commands for Interfaces

clear ipv6 neighbors

Syntax	clear ipv6 neighbors <all host <i>hostname</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.3 for EX Series switches.
Description	Clear IPv6 neighbor cache information.
Options	none—Clear all IPv6 neighbor cache information. all—(Optional) Clear all IPv6 neighbor cache information. host <i>hostname</i> —(Optional) Clear the information for the specified IPv6 neighbors.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• show ipv6 neighbors on page 277
List of Sample Output	clear ipv6 neighbors on page 202
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear ipv6 neighbors user@host> clear ipv6 neighbors

monitor interface

Syntax	monitor interface <interface-name traffic <detail>>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display real-time statistics about interfaces, updating the statistics every second. Check for and display common interface failures, such as SONET/SDH and T3 alarms, loopbacks detected, and increases in framing errors.
Options	<p>none—Display real-time statistics for all interfaces.</p> <p>interface-name—(Optional) Display real-time statistics for the specified interface.</p> <p>traffic—(Optional) Display traffic data for all active interfaces.</p> <p>detail—(Optional) With traffic option only, display detailed output.</p>
Additional Information	The output of this command shows how much each field has changed since you started the command or since you cleared the counters by using the c key. For a description of the statistical information provided in the output of this command, see the show interfaces extensive command for a particular interface type in the Junos OS Interfaces Command Reference . To control the output of the monitor interface interface-name command while it is running, use the keys listed in Table 28 on page 203 . The keys are not case-sensitive.

Table 28: Output Control Keys for the monitor interface interface-name Command

Key	Action
c	Clears (returns to zero) the delta counters since monitor interface was started. This does not clear the accumulative counter. To clear the accumulative counter, use the clear interfaces interval command.
f	Freezes the display, halting the display of updated statistics and delta counters.
i	Displays information about a different interface. The command prompts you for the name of a specific interface.
n	Displays information about the next interface. The monitor interface command displays the physical or logical interfaces in the same order as the show interfaces terse command.
q or Esc	Quits the command and returns to the command prompt.
t	Thaws the display, resuming the update of the statistics and delta counters.

To control the output of the **monitor interface traffic** command while it is running, use the keys listed in [Table 29 on page 204](#). The keys are not case-sensitive.

Table 29: Output Control Keys for the monitor interface traffic Command

Key	Action
b	Displays the statistics in units of bytes and bytes per second (bps).
c	Clears (return to 0) the delta counters in the Current Delta column. The statistics counters are not cleared.
d	Displays the Current Delta column (instead of the rate column) in bps or packets per second (pps).
p	Displays the statistics in units of packets and packets per second (pps).
q or Esc	Quits the command and returns to the command prompt.
r	Displays the rate column (instead of the Current Delta column) in bps and pps.

Required Privilege Level trace

List of Sample Output [monitor interface \(Physical\) on page 205](#)
[monitor interface \(OTN Interface\) on page 207](#)
[monitor interface \(Logical\) on page 208](#)
[monitor interface traffic on page 208](#)
[monitor interface traffic detail on page 209](#)

Output Fields [Table 30 on page 204](#) describes the output fields for the **monitor interface** command. Output fields are listed in the approximate order in which they appear.

Table 30: monitor interface Output Fields

Field Name	Field Description	Level of Output
routerl	Hostname of the router.	All levels
Seconds	How long the monitor interface command has been running or how long since you last cleared the counters.	All levels
Time	Current time (UTC).	All levels
Delay x/y/z	Time difference between when the statistics were displayed and the actual clock time. <ul style="list-style-type: none"> x—Time taken for the last polling (in milliseconds). y—Minimum time taken across all pollings (in milliseconds). z—Maximum time taken across all pollings (in milliseconds). 	All levels
Interface	Short description of the interface, including its name, status, and encapsulation.	All levels
Link	State of the link: Up , Down , or Test .	All levels

Table 30: monitor interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
Current delta	Cumulative number for the counter in question since the time shown in the Seconds field, which is the time since you started the command or last cleared the counters.	All levels
Local Statistics	<p>(Logical interfaces only) Number and rate of bytes and packets destined to the router or switch through the specified interface. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize.:</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. 	All levels
Remote Statistics	<p>(Logical interfaces only) Statistics for traffic transiting the router or switch. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize.:</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. 	All levels
Traffic statistics	<p>Total number of bytes and packets received and transmitted on the interface. These statistics are the sum of the local and remote statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. 	All levels
Description	With the traffic option, displays the interface description configured at the [edit interfaces <i>interface-name</i>] hierarchy level.	detail

Sample Output

```

monitor interface user@host> monitor interface so-0/0/0
(Physical)         router1                               Seconds: 19           Time: 15:46:29

Interface: so-0/0/0, Enabled, Link is Up
Encapsulation: PPP, Keepalives, Speed: 0C48
Traffic statistics:
Input packets:           6045 (0 pps)
Input bytes:             6290065 (0 bps)
Output packets:         10376 (0 pps)
Output bytes:           10365540 (0 bps)

Current Delta
[11]
[13882]
[10]
[9418]
```

```
Encapsulation statistics:
  Input keepalives:          1901          [2]
  Output keepalives:        1901          [2]
  NCP state: Opened
  LCP state: Opened
Error statistics:
  Input errors:              0             [0]
  Input drops:               0             [0]
  Input framing errors:      0             [0]
  Policed discards:          0             [0]
  L3 incompletes:            0             [0]
  L2 channel errors:         0             [0]
  L2 mismatch timeouts:      0             [0]
  Carrier transitions:        1             [0]
  Output errors:             0             [0]
  Output drops:              0             [0]
  Aged packets:              0             [0]
Active alarms : None
Active defects: None
SONET error counts/seconds:
  LOS count                  1             [0]
  LOF count                   1             [0]
  SEF count                   1             [0]
  ES-S                        0             [0]
  SES-S                       0             [0]
SONET statistics:
  BIP-B1                     458871         [0]
  BIP-B2                     460072         [0]
  REI-L                       465610         [0]
  BIP-B3                     458978         [0]
  REI-P                       458773         [0]
```

```

Received SONET overhead:
  F1      : 0x00  J0      : 0x00  K1      : 0x00
  K2      : 0x00  S1      : 0x00  C2      : 0x00
  C2(cmp) : 0x00  F2      : 0x00  Z3      : 0x00
  Z4      : 0x00  S1(cmp) : 0x00
Transmitted SONET overhead:
  F1      : 0x00  J0      : 0x01  K1      : 0x00
  K2      : 0x00  S1      : 0x00  C2      : 0xcf
  F2      : 0x00  Z3      : 0x00  Z4      : 0x00

```

Next='n', Quit='q' or ESC, Freeze='f', Thaw='t', Clear='c', Interface='i'

monitor interface (OTN Interface)

```
user@host> monitor interface ge-7/0/0
```

```

Interface: ge-7/0/0, Enabled, Link is Up
Encapsulation: Ethernet, Speed: 10000Mbps
Traffic statistics:
  Input bytes:                0 (0 bps)
  Output bytes:               0 (0 bps)
  Input packets:              0 (0 pps)
  Output packets:             0 (0 pps)
Error statistics:
  Input errors:                0
  Input drops:                 0
  Input framing errors:        0
  Policed discards:            0
  L3 incompletes:              0
  L2 channel errors:           0
  L2 mismatch timeouts:        0
  Carrier transitions:         5
  Output errors:               0
  Output drops:                0
  Aged packets:                0
Active alarms : None
Active defects: None
Input MAC/Filter statistics:
  Unicast packets              0
  Broadcast packets            0
  Multicast packets            0
  Oversized frames             0
  Packet reject count          0
  DA rejects                   0
  SA rejects                   0
Output MAC/Filter Statistics:
  Unicast packets              0
  Broadcast packets            0
  Multicast packets            0
  Packet pad count             0
  Packet error count           0
OTN Link 0
  OTN Alarms: OTU_BDI, OTU_TTIM, ODU_BDI
  OTN Defects: OTU_BDI, OTU_TTIM, ODU_BDI, ODU_TTIM
  OTN OC - Seconds
    LOS                        2
    LOF                        9
  OTN OTU - FEC Statistics
    Corr err ratio             N/A
    Corr bytes                  0
    Uncorr words                0
  OTN OTU - Counters

```

```

BIP                                0
BBE                                0
ES                                 0
SES                                0
UAS                                422
OTN ODU - Counters
BIP                                0
BBE                                0
ES                                 0
SES                                0
UAS                                422
OTN ODU - Received Overhead    APSGCC 0-3:          0

```

```

monitor interface user@host> monitor interface so-1/0/0.0
(Logical)          host name                Seconds: 16                Time: 15:33:39
                                                Delay: 0/0/1

Interface: so-1/0/0.0, Enabled, Link is Down
Flags: Hardware-Down Point-To-Point SNMP-Traps
Encapsulation: PPP
Local statistics:
Input bytes:                0                Current delta [0]
Output bytes:               0                [0]
Input packets:              0                [0]
Output packets:             0                [0]
Remote statistics:
Input bytes:                0 (0 bps)         [0]
Output bytes:               0 (0 bps)         [0]
Input packets:              0 (0 pps)         [0]
Output packets:             0 (0 pps)         [0]
Traffic statistics:
Destination address: 192.168.8.193, Local: 192.168.8.21

Next='n', Quit='q' or ESC, Freeze='f', Thaw='t', Clear='c', Interface='i'

```

```

monitor interface user@host> monitor interface traffic
traffic          host name                Seconds: 15                Time: 12:31:09

Interface  Link  Input packets  (pps)  Output packets  (pps)
so-1/0/0   Down    0              (0)    0              (0)
so-1/1/0   Down    0              (0)    0              (0)
so-1/1/1   Down    0              (0)    0              (0)
so-1/1/2   Down    0              (0)    0              (0)
so-1/1/3   Down    0              (0)    0              (0)
t3-1/2/0   Down    0              (0)    0              (0)
t3-1/2/1   Down    0              (0)    0              (0)
t3-1/2/2   Down    0              (0)    0              (0)
t3-1/2/3   Down    0              (0)    0              (0)
so-2/0/0   Up      211035         (1)    36778          (0)
so-2/0/1   Up      192753         (1)    36782          (0)
so-2/0/2   Up      211020         (1)    36779          (0)
so-2/0/3   Up      211029         (1)    36776          (0)
so-2/1/0   Up      189378         (1)    36349          (0)
so-2/1/1   Down    0              (0)    18747          (0)
so-2/1/2   Down    0              (0)    16078          (0)
so-2/1/3   Up      0              (0)    80338          (0)
at-2/3/0   Up      0              (0)    0              (0)
at-2/3/1   Down    0              (0)    0              (0)

Bytes=b, Clear=c, Delta=d, Packets=p, Quit=q or ESC, Rate=r, Up=^U, Down=^D

```

```
monitor interface user@host> monitor interface traffic detail
traffic detail host name Seconds: 15 Time: 12:31:09

Interface Link Input packets (pps) Output packets (pps) Description
t1-0/1/1:0 Up 19769 (0) 0 (0) To-OSAKA-1
...
Bytes=b, Clear=c, Delta=d, Packets=p, Quit=q or ESC, Rate=r, Up=^U, Down=^D
```

request diagnostics tdr

Syntax request diagnostics tdr (abort | start) interface *interface-name*

Release Information Command introduced in Junos OS Release 9.0 for EX Series switches.

Description Start a time domain reflectometry (TDR) diagnostic test on the specified interface. This test characterizes and locates faults on twisted-pair Ethernet cables. For example, it can detect a broken twisted pair and provide the approximate distance to the break. It can also detect polarity swaps, pair swaps, and excessive skew.

The TDR test is supported on the following switches and interfaces:

- EX2200, EX3200, EX3300, and EX4200 switches—RJ-45 network interfaces. The TDR test is not supported on management interfaces and SFP interfaces.
- EX6200 and EX8200 switches—RJ-45 interfaces on line cards.



NOTE: We recommend running the TDR test when there is no traffic on the interface under test.

You view the results of the TDR test with the [show diagnostics tdr](#) command.

Options **abort**—Stop the TDR test currently in progress on the specified interface. No results are reported, and previous results, if any, are cleared.

interface-name—The name of the interface.

start—Start a TDR test on the specified interface.

Required Privilege Level maintenance

Related Documentation

- [show diagnostics tdr on page 212](#)
- [Diagnosing a Faulty Twisted-Pair Cable \(CLI Procedure\) on page 126](#)

List of Sample Output [request diagnostics tdr start interface ge-0/0/19 on page 211](#)

Output Fields [Table 31 on page 211](#) lists the output fields for the **request diagnostics tdr** command. Output fields are listed in the approximate order in which they appear.

Table 31: 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      user@switch> request diagnostics tdr start interface ge-0/0/19
start interface
ge-0/0/19                  Interface TDR detail:
                           Test status                : Test successfully executed ge-0/0/19
```

show diagnostics tdr

Syntax	<code>show diagnostics tdr</code> <code><interface <i>interface-name</i>></code>
Release Information	Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	<p>Display the results of a time domain reflectometry (TDR) diagnostic test run on an interface. A TDR test characterizes and locates faults on twisted-pair Ethernet cables. For example, it can detect a broken twisted pair and provide the approximate distance to the break. It can also detect polarity swaps, pair swaps, and excessive skew.</p> <p>The TDR test is supported on the following switches and interfaces:</p> <ul style="list-style-type: none">• EX2200, EX3200, EX3300, and EX4200 switches—RJ-45 network interfaces. The TDR test is not supported on management interfaces and SFP interfaces.• EX6200 and EX8200 switches— RJ-45 interfaces on line cards. <p>Use the request diagnostics tdr command to request a TDR test on a specified interface. Use the show diagnostic tdr command to display the last TDR test results for a specified interface or the last TDR test results for all network interfaces on the switch that support the TDR test.</p>
Options	<p><code>none</code>—Show summarized last results for all interfaces on the switch that support the TDR test.</p> <p><code>interface <i>interface-name</i></code>—(Optional) Show detailed last results for the specified interface or a range of interfaces. Specify a range of interfaces by entering the beginning and ending interface in the range, separated by a dash—for example, <code>ge-0/0/15-ge-0/0/20</code>.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• request diagnostics tdr on page 210• Diagnosing a Faulty Twisted-Pair Cable (CLI Procedure) on page 126
List of Sample Output	show diagnostics tdr interface ge-0/0/19 (Normal Cable) on page 214 show diagnostics tdr interface ge-2/0/2 (Faulty Cable) on page 215 show diagnostics tdr (All Supported Interfaces) on page 215
Output Fields	Table 32 on page 213 lists the output fields for the show diagnostics tdr command. Output fields are listed in the approximate order in which they appear.

Table 32: show diagnostics tdr Output Fields

Field Name	Field Description
Interface name or Interface	Name of interface for which TDR test results are being reported.
Test status	<p>Status of TDR test:</p> <ul style="list-style-type: none"> • Aborted—Test was terminated by operator before it was complete. • Failed—Test was not completed successfully. • Interface <i>interface-name</i> not found—Specified interface does not exist. • Not Started—No TDR test results are available for the interface. • Passed—Test completed successfully. The cable, however, might still have a fault—see the Cable status field for information on the cable. • Started—Test is currently running and not yet complete. • VCT not supported on <i>interface-name</i>—TDR test is not supported on the interface.
Link status	Operating status of link: UP or Down .
MDI pair	Twisted pair for which test results are being reported, identified by pin numbers. (Displayed only when the interface option is used.)
Cable status	<p>When detailed information is displayed, status for a twisted pair:</p> <ul style="list-style-type: none"> • Failed—TDR test failed on the cable pair. • Impedance Mismatch—Impedance on the twisted pair is not correct. Possible reasons for an impedance mismatch include: <ul style="list-style-type: none"> • The twisted pair is not connected properly. • The twisted pair is damaged. • The connector is faulty. • Normal—No cable fault detected for the twisted pair. • Open—Lack of continuity between the pins at each end of the twisted-pair. • Short on Pair-<i>n</i>—A short-circuit was detected on the twisted pair. <p>When summary information for all interfaces is displayed, status for the cable as a whole:</p> <ul style="list-style-type: none"> • Fault—A fault was detected on one or more of the twisted-pairs. • OK—No fault was detected on any of the twisted pairs.
Distance fault or Max distance fault	<p>Distance to the fault in whole meters. If there is no fault, this value is 0.</p> <p>When summary information for all interfaces is displayed, this value is the distance to the most distant fault if there is more than one twisted pair with a fault.</p>

Table 32: show diagnostics tdr Output Fields (*continued*)

Field Name	Field Description
Polarity swap	<p>Indicates the polarity status of the twisted pair:</p> <ul style="list-style-type: none"> • Normal—Polarity is normal. Each conductor in the twisted pair has been connected the same pins at the both ends of the connection. For example, a conductor connected to pin 1 at the near end of the connection is connected to pin 1 at the far end. • Reversed—Polarity has been reversed. For the twisted pair, the conductors have switched which pins they are connected to at the near and far ends of the connection. For example, the conductor connected to pin 1 at the near end is connected to pin 2 at the far end. <p>(Not available on EX8200 switches.) (Displayed only when the interface option is used)</p>
Skew time	<p>Difference in nanoseconds between the propagation delay on this twisted pair and the twisted pair with the shortest propagation delay. (Not available on EX8200 switches.) (Displayed only when the interface option is used.)</p>
Channel Pair	<p>Number of the 10/100BASE-T transmit/receive pair being reported on.</p>
Pair Swap	<p>Indicates whether or not the twisted pairs are swapped:</p> <ul style="list-style-type: none"> • MDI—The pairs are not swapped (straight-through cable). • MDIX—The pairs are swapped (cross-over cable). <p>(Displayed only when the interface option is used.)</p>
Downshift	<p>Indicates whether the connection speed is being downshifted:</p> <ul style="list-style-type: none"> • No Downshift—No downshifting of connection speed. • Downshift occurs—Connection speed is downshifted to 10 or 100 Mbs. This occurs if the cable is a two-pair cable rather than the four-pair cable required by Gigabit Ethernet. <p>(Displayed only when the interface option is used.)</p>

Sample Output

```

show diagnostics tdr interface ge-0/0/19
interface ge-0/0/19
  (Normal Cable)
user@switch> show diagnostics tdr interface ge-0/0/19
Interface TDR detail:
Interface name      : ge-0/0/19
Test status         : Passed
Link status         : UP
MDI pair            : 1-2
  Cable status      : Normal
  Distance fault    : 0 Meters
  Polartiy swap     : Normal
  Skew time         : 0 ns
MDI pair            : 3-6
  Cable status      : Normal
  Distance fault    : 0 Meters

```

```

Polartiy swap           : Normal
Skew time               : 8 ns
MDI pair                : 4-5
Cable status            : Normal
Distance fault          : 0 Meters
Polartiy swap           : Normal
Skew time               : 8 ns
MDI pair                : 7-8
Cable status            : Normal
Distance fault          : 0 Meters
Polartiy swap           : Normal
Skew time               : 8 ns
Channel pair            : 1
Pair swap               : MDI
Channel pair            : 2
Pair swap               : MDI
Downshift               : No Downshift

```

show diagnostics tdr interface ge-2/0/2 (Faulty Cable)

```

user@switch> show diagnostics tdr interface ge-2/0/2
Interface TDR detail:
Interface name          : ge-2/0/2
Test status             : Passed
Link status             : Down
MDI Pair               : 1-2
  Cable status          : 1-2
  Distance fault        : 2 Meters
  Polartiy swap         : N/A
  Skew time             : N/A
MDI Pair               : 3-6
  Cable status          : Impedance Mismatch
  Distance fault        : 3 Meters
  Polartiy swap         : N/A
  Skew time             : N/A
MDI Pair               : 4-5
  Cable status          : Impedance Mismatch
  Distance fault        : 3 Meters
  Polartiy swap         : N/A
  Skew time             : N/A
MDI Pair               : 7-8
  Cable status          : Short on Pair-2
  Distance fault        : 3 Meters
  Polartiy swap         : N/A
  Skew time             : N/A
Channel pair           : 1
Pair swap              : N/A
Channel pair           : 2
Pair swap              : N/A
Downshift              : N/A

```

show diagnostics tdr (All Supported Interfaces)

```

user@switch> show diagnostics tdr

```

Interface	Test status	Link status	Cable status	Max distance fault
ge-0/0/0	Not Started	N/A	N/A	N/A
ge-0/0/1	Not Started	N/A	N/A	N/A
ge-0/0/2	Started	N/A	N/A	N/A
ge-0/0/3	Started	N/A	N/A	N/A
ge-0/0/4	Passed	UP	OK	0
ge-0/0/5	Passed	UP	Fault	173
ge-0/0/6	Passed	UP	OK	0
ge-0/0/7	Passed	UP	OK	0
ge-0/0/8	Passed	UP	OK	0
ge-0/0/9	Passed	UP	OK	0

ge-0/0/10	Passed	UP	OK	0
ge-0/0/11	Passed	UP	OK	0
ge-0/0/12	Passed	UP	OK	0
ge-0/0/13	Passed	UP	OK	0
ge-0/0/14	Passed	UP	OK	0
ge-0/0/15	Passed	UP	OK	0
ge-0/0/16	Passed	UP	OK	0
ge-0/0/17	Passed	UP	OK	0
ge-0/0/18	Passed	UP	OK	0
ge-0/0/19	Passed	UP	OK	0
ge-0/0/20	Passed	Down	Fault	0
ge-0/0/21	Passed	Down	Fault	5
ge-0/0/22	Passed	UP	OK	0
ge-0/0/23	Passed	UP	OK	0

show ethernet-switching interfaces

Syntax	<pre>show ethernet-switching interfaces <brief detail summary> <interface <i>interface-name</i>></pre>
Release Information	<p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>In Junos OS Release 9.6 for EX Series switches, the following updates were made:</p> <ul style="list-style-type: none"> • Blocking field output was updated. • The default view was updated to include information about 802.1Q tags. • The detail view was updated to include information on VLAN mapping. <p>In Junos OS Release 11.1 for EX Series switches, the detail view was updated to include reflective relay information.</p>
Description	Display information about Ethernet switching interfaces.
Options	<p>none—Display brief information for Ethernet switching interfaces.</p> <p>brief detail summary—(Optional) Display the specified level of output.</p> <p>interface <i>interface-name</i>—(Optional) Display Ethernet switching information for a specific interface.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show ethernet-switching mac-learning-log • show ethernet-switching table • Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces (CLI Procedure)
List of Sample Output	<p>show ethernet-switching interfaces on page 219</p> <p>show ethernet-switching interfaces ge-0/0/15 brief on page 219</p> <p>show ethernet-switching interfaces ge-0/0/2 detail (Blocked by RTG rtggroup) on page 219</p> <p>show ethernet-switching interfaces ge-0/0/15 detail (Blocked by STP) on page 220</p> <p>show ethernet-switching interfaces ge-0/0/17 detail (Disabled by bpdu-control) on page 220</p> <p>show ethernet-switching interfaces detail (C-VLAN to S-VLAN Mapping) on page 220</p> <p>show ethernet-switching interfaces detail (Reflective Relay Is Configured) on page 220</p>
Output Fields	<p>Table 33 on page 218 lists the output fields for the show ethernet-switching interfaces command. Output fields are listed in the approximate order in which they appear.</p>

Table 33: show ethernet-switching interfaces Output Fields

Field Name	Field Description	Level of Output
Interface	Name of a switching interface.	none, brief , detail , summary
Index	VLAN index internal to Junos OS.	detail
State	Interface state. Values are up and down .	none, brief , detail
Port mode	The access mode is the port mode default and works with a single VLAN. Port mode can also be trunk , which accepts tagged packets from multiple VLANs on other switches. The third port mode value is tagged-access , which accepts tagged packets from access devices.	detail
Reflective Relay Status	Reflective relay allows packets to use the same interface for both upstream and downstream traffic. When reflective relay has been configured, the status displayed is always enabled . When reflective relay is not configured, this entry does not appear in the command output.	detail
Ether type for the interface	Ether type is a two-octet field in an Ethernet frame used to indicate which protocol is encapsulated in the payload of an incoming Ethernet packet. Both 802.1Q packets and Q-in-Q packets use this field. The output displayed for this particular field indicates the interface's Ether type, which is used to match the Ether type of incoming 802.1Q packets and Q-in-Q packets. The indicated Ether type field is also added to the interface's outgoing 802.1Q and Q-in-Q packets.	detail
VLAN membership	Names of VLANs that belong to this interface.	none, brief , detail ,
Tag	Number of the 802.1Q tag.	none, brief , detail ,
Tagging	Specifies whether the interface forwards 802.1Q tagged or untagged traffic.	none, brief , detail ,
Blocking	<p>The forwarding state of the interface:</p> <ul style="list-style-type: none"> • unblocked—Traffic is forwarded on the interface. • blocked—Traffic is not being forwarded on the interface. • Disabled by bpdu control—The interface is disabled due to receiving BPDUs on a protected interface. If the disable-timeout statement has been included in the BPDU configuration, the interface automatically returns to service after the timer expires. • blocked by RTG—The specified redundant trunk group is disabled. • blocked by STP—The interface is disabled due to a spanning-tree protocol error. • MAC limit exceeded—The interface is temporarily disabled due to a MAC limit error. The disabled interface is automatically restored to service when the disable timeout expires. • MAC move limit exceeded—The interface is temporarily disabled due to a MAC move limit error. The disabled interface is automatically restored to service when the disable timeout expires. • Storm control in effect—The interface is temporarily disabled due to a storm control error. The disabled interface is automatically restored to service when the disable timeout expires. 	none, brief , detail ,

Table 33: show ethernet-switching interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
Number of MACs learned on IFL	Number of MAC addresses learned by this interface.	detail
mapping	<p>When mapping is configured, the status is one of the following C-VLAN to S-VLAN mapping types:</p> <ul style="list-style-type: none"> • dot1q-tunneled—The interface maps all traffic to the S-VLAN (all-in-one bundling). • native—The interface maps untagged and priority tagged packets to the S-VLAN. • push—The interface maps packets to a firewall filter to an S-VLAN. • policy-mapped—The interface maps packets to a specifically defined S-VLAN. • integer—The interface maps packets to the specified S-VLAN. <p>When mapping is not configured, this entry does not appear in the command output.</p>	detail

Sample Output

```

show user@switch> show ethernet-switching interfaces
ethernet-switching
interfaces
Interface    State  VLAN members    Tag  Tagging  Blocking
-----
ae0.0        up     default
ge-0/0/2.0   up     vlan300         300   untagged  unblocked
ge-0/0/3.0   up     default         blocked by STP
ge-0/0/4.0   down   default         MAC limit exceeded
ge-0/0/5.0   down   default         MAC move limit exceeded
ge-0/0/6.0   down   default         Storm control in effect
ge-0/0/7.0   down   default         unblocked
ge-0/0/13.0  up     default         untagged  unblocked
ge-0/0/14.0  up     vlan100         100   tagged   unblocked
              vlan200         200   tagged   unblocked
ge-0/0/15.0  up     vlan100         100   tagged   blocked by STP
              vlan200         200   tagged   blocked by STP
ge-0/0/16.0  down   default         untagged  unblocked
ge-0/0/17.0  down   vlan100         100   tagged   Disabled by bpdu-control
              vlan200         200   tagged   Disabled by bpdu-control

show user@switch> show ethernet-switching interfaces ge-0/0/15 brief
ethernet-switching
interfaces ge-0/0/15
brief
Interface    State  VLAN members    Tag  Tagging  Blocking
-----
ge-0/0/15.0  up     vlan100         100   tagged   blocked by STP
              vlan200         200   tagged   blocked by STP

show user@switch> show ethernet-switching interfaces ge-0/0/2 detail
ethernet-switching
interfaces ge-0/0/2
detail (Blocked by RTG
rtggroup)
Interface: ge-0/0/2.0, Index: 65, State: up, Port mode: Access
Ether type for the interface: 0X8100
VLAN membership:
vlan300, 802.1Q Tag: 300, untagged, msti-id: 0, blocked by RTG(rtgroup)

```

Number of MACs learned on IFL: 0

**show
ethernet-switching
interfaces ge-0/0/15
detail (Blocked by
STP)**

user@switch> show ethernet-switching interfaces ge-0/0/15 detail

Interface: ge-0/0/15.0, Index: 70, State: up, Port mode: Trunk
Ether type for the interface: 0X8100
VLAN membership:
 vlan100, 802.1Q Tag: 100, tagged, msti-id: 0, blocked by STP
 vlan200, 802.1Q Tag: 200, tagged, msti-id: 0, blocked by STP

Number of MACs learned on IFL: 0

**show
ethernet-switching
interfaces ge-0/0/17
detail (Disabled by
bpdu-control)**

user@switch> show ethernet-switching interfaces ge-0/0/17 detail

Interface: ge-0/0/17.0, Index: 71, State: down, Port mode: Trunk
Ether type for the interface: 0X8100
VLAN membership:
 vlan100, 802.1Q Tag: 100, tagged, msti-id: 1, Disabled by bpdu-control
 vlan200, 802.1Q Tag: 200, tagged, msti-id: 2, Disabled by bpdu-control
Number of MACs learned on IFL: 0

**show
ethernet-switching
interfaces detail
(C-VLAN to S-VLAN
Mapping)**

user@switch> show ethernet-switching interfaces ge-0/0/6.0 detail

Interface: ge-0/0/6.0, Index: 73, State: up, Port mode: Access
Ether type for the interface: 0X8100
VLAN membership:
 map, 802.1Q Tag: 134, Mapped Tag: native, push, dot1q-tunneled, unblocked
 map, 802.1Q Tag: 134, Mapped Tag: 20, push, dot1q-tunneled, unblocked

**show
ethernet-switching
interfaces detail
(Reflective Relay Is
Configured)**

user@switch1> show ethernet-switching interfaces ge-7/0/2 detail

Interface: ge-7/0/2, Index: 66, State: down, Port mode: Tagged-access
Ether type for the interface: 0X8100
Reflective Relay Status: Enabled
Ether type for the interface: 0x8100
VLAN membership:
 VLAN_Purple VLAN_Orange VLAN_Blue, 802.1Q Tag: 450, tagged, unblocked
Number of MACs learned on IFL: 0

show interfaces diagnostics optics

Syntax	<code>show interfaces diagnostics optics <i>interface-name</i></code>
Release Information	Command introduced in Junos OS Release 10.0 for EX Series switches.
Description	<p>Display diagnostics data and alarms for Gigabit Ethernet optical transceivers (SFP, SFP+, or XFP) installed in 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 the optics module is not operating properly. This information can be used to diagnose why a transceiver is not working.</p>
Options	<i>interface-name</i> —Name of the interface associated with the port in which the transceiver is installed: ge-<i>fpc/pic/port</i> or xe-<i>fpc/pic/port</i> .
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Monitoring Interface Status and Traffic on page 113 • Installing a Transceiver in an EX Series Switch • Removing a Transceiver from an EX Series Switch • Junos OS Ethernet Interfaces Configuration Guide
List of Sample Output	show interfaces diagnostics optics ge-0/1/0 (SFP Transceiver) on page 225 show interfaces diagnostics optics xe-0/1/0 (SFP+ Transceiver) on page 226 show interfaces diagnostics optics xe-0/1/0 (XFP Transceiver) on page 226
Output Fields	Table 34 on page 221 lists the output fields for the show interfaces diagnostics optics command. Output fields are listed in the approximate order in which they appear.

Table 34: show interfaces diagnostics optics Output Fields

Field Name	Field Description
Physical interface	Displays the name of the physical interface.
Laser bias current	Displays the magnitude of the laser bias power setting current, in milliamperes. 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.

Table 34: show interfaces diagnostics optics Output Fields (*continued*)

Field Name	Field Description
Module voltage (Not available for XFP transceivers)	Displays the voltage, in Volts.
Laser rx power (Not available for SFP and SFP+ transceivers)	Displays the laser received optical power, in milliwatts (mW) and decibels referred to 1.0 mW (dBm).
Receiver signal average optical power (Not available for XFP transceivers)	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 On or Off .
Laser bias current low alarm	Displays whether the laser bias power setting low alarm is On or Off .
Laser bias current high warning	Displays whether the laser bias power setting high warning is On or Off .
Laser bias current low warning	Displays whether the laser bias power setting low warning is On or Off .
Laser output power high alarm	Displays whether the laser output power high alarm is On or Off .
Laser output power low alarm	Displays whether the laser output power low alarm is On or Off .
Laser output power high warning	Displays whether the laser output power high warning is On or Off .
Laser output power low warning	Displays whether the laser output power low warning is On or Off .
Module temperature high alarm	Displays whether the module temperature high alarm is On or Off .
Module temperature low alarm	Displays whether the module temperature low alarm is On or Off .
Module temperature high warning	Displays whether the module temperature high warning is On or Off .
Module temperature low warning	Displays whether the module temperature low warning is On or Off .
Module voltage high alarm (Not available for XFP transceivers)	Displays whether the module voltage high alarm is On or Off .
Module voltage low alarm (Not available for XFP transceivers)	Displays whether the module voltage low alarm is On or Off .
Module voltage high warning (Not available for XFP transceivers)	Displays whether the module voltage high warning is On or Off .

Table 34: show interfaces diagnostics optics Output Fields (*continued*)

Field Name	Field Description
Module voltage low warning (Not available for XFP transceivers)	Displays whether the module voltage low warning is On or Off .
Laser rx power high alarm	Displays whether the receive laser power high alarm is On or Off .
Laser rx power low alarm	Displays whether the receive laser power low alarm is On or Off .
Laser rx power high warning	Displays whether the receive laser power high warning is On or Off .
Laser rx power low warning	Displays whether the receive laser power low warning is On or Off .
Laser bias current high alarm threshold	Displays the vendor-specified threshold for the laser bias current high alarm.
Module not ready alarm (Not available for SFP and SFP+ transceivers)	Displays whether the module not ready alarm is On or Off . When the output is On , the module has an operational fault.
Module power down alarm (Not available for SFP and SFP+ transceivers)	Displays whether the module power down alarm is On or Off . When the output is On , module is in a limited power mode, low for normal operation.
Tx data not ready alarm (Not available for SFP and SFP+ transceivers)	Any condition leading to invalid data on the transmit path. Displays whether the Tx data not ready alarm is On or Off .
Tx not ready alarm (Not available for SFP and SFP+ transceivers)	Any condition leading to invalid data on the transmit path. Displays whether the Tx not ready alarm is On or Off .
Tx laser fault alarm (Not available for SFP and SFP+ transceivers)	Laser fault condition. Displays whether the Tx laser fault alarm is On or Off .
Tx CDR loss of lock alarm (Not available for SFP and SFP+ transceivers)	Transmit clock and data recovery (CDR) loss of lock. Loss of lock on the transmit side of the CDR. Displays whether the Tx CDR loss of lock alarm is On or Off .
Rx not ready alarm (Not available for SFP and SFP+ transceivers)	Any condition leading to invalid data on the receive path. Displays whether the Rx not ready alarm is On or Off .
Rx loss of signal alarm (Not available for SFP and SFP+ transceivers)	Receive loss of signal alarm. When on , indicates insufficient optical input power to the module. Displays whether the Rx loss of signal alarm is On or Off .
Rx CDR loss of lock alarm (Not available for SFP and SFP+ transceivers)	Receive CDR loss of lock. Loss of lock on the receive side of the CDR. Displays whether the Rx CDR loss of lock alarm is On or Off .

Table 34: show interfaces diagnostics optics Output Fields (*continued*)

Field Name	Field Description
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.
Module voltage high alarm threshold (Not available for XFP transceivers)	Displays the vendor-specified threshold for the module voltage high alarm.
Module voltage low alarm threshold (Not available for XFP transceivers)	Displays the vendor-specified threshold for the module voltage low alarm.
Module voltage high warning threshold (Not available for XFP transceivers)	Displays the vendor-specified threshold for the module voltage high warning.
Module voltage low warning threshold (Not available for XFP transceivers)	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.

Table 34: show interfaces diagnostics optics Output Fields (*continued*)

Field Name	Field Description
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 interfaces user@host> show interfaces diagnostics optics ge-0/1/0
diagnostics optics Physical interface: ge-0/1/0
ge-0/1/0          Laser bias current           : 5.444 mA
(SFP Transceiver) Laser output power           : 0.3130 mW / -5.04 dBm
                  Module temperature       : 36 degrees C / 97 degrees F
                  Module voltage           : 3.2120 V
                  Receiver signal average optical power : 0.3840 mW / -4.16 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 : 15.000 mA
                  Laser bias current low alarm threshold : 1.000 mA
                  Laser bias current high warning threshold : 12.000 mA
                  Laser bias current low warning threshold : 2.000 mA
                  Laser output power high alarm threshold : 0.6300 mW / -2.01 dBm
                  Laser output power low alarm threshold : 0.0660 mW / -11.80 dBm
                  Laser output power high warning threshold : 0.6300 mW / -2.01 dBm
                  Laser output power low warning threshold : 0.0780 mW / -11.08 dBm
                  Module temperature high alarm threshold : 109 degrees C / 228 degrees F
                  Module temperature low alarm threshold : -29 degrees C / -20 degrees F
                  Module temperature high warning threshold : 103 degrees C / 217 degrees F
                  Module temperature low warning threshold : -13 degrees C / 9 degrees F
                  Module voltage high alarm threshold : 3.900 V
                  Module voltage low alarm threshold : 2.700 V
                  Module voltage high warning threshold : 3.700 V
                  Module voltage low warning threshold : 2.900 V
                  Laser rx power high alarm threshold : 1.2589 mW / 1.00 dBm
                  Laser rx power low alarm threshold : 0.0100 mW / -20.00 dBm

```

```
Laser rx power high warning threshold : 0.7939 mW / -1.00 dBm
Laser rx power low warning threshold  : 0.0157 mW / -18.04 dBm
```

Sample Output

```
show interfaces user@host> show interfaces diagnostics optics xe-0/1/0
diagnostics optics Physical interface: xe-0/1/0
xe-0/1/0          Laser bias current           : 4.968 mA
(SFP+ Transceiver) Laser output power          : 0.4940 mW / -3.06 dBm
                  Module temperature       : 27 degrees C / 81 degrees F
                  Module voltage          : 3.2310 V
                  Receiver signal average optical power : 0.0000
                  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               : On
                  Laser rx power high warning            : Off
                  Laser rx power low warning             : On
                  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
```

Sample Output

```
show interfaces user@host> show interfaces diagnostics optics xe-0/1/0
diagnostics optics Physical interface: xe-0/1/0
xe-0/1/0          Laser bias current           : 8.029 mA
(XFP Transceiver) Laser output power          : 0.6430 mW / -1.92 dBm
                  Module temperature       : 4 degrees C / 39 degrees F
```



```

Laser rx power : 0.0012 mW / -29.21 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
Laser rx power high alarm : Off
Laser rx power low alarm : On
Laser rx power high warning : Off
Laser rx power low warning : On
Module not ready alarm : On
Module power down alarm : Off
Tx data not ready alarm : Off
Tx not ready alarm : Off
Tx laser fault alarm : Off
Tx CDR loss of lock alarm : Off
Rx not ready alarm : On
Rx loss of signal alarm : On
Rx CDR loss of lock alarm : On
Laser bias current high alarm threshold : 13.000 mA
Laser bias current low alarm threshold : 2.000 mA
Laser bias current high warning threshold : 12.000 mA
Laser bias current low warning threshold : 3.000 mA
Laser output power high alarm threshold : 0.8310 mW / -0.80 dBm
Laser output power low alarm threshold : 0.1650 mW / -7.83 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 : 90 degrees C / 194 degrees F
Module temperature low alarm threshold : 0 degrees C / 32 degrees F
Module temperature high warning threshold : 85 degrees C / 185 degrees F
Module temperature low warning threshold : 0 degrees C / 32 degrees F
Laser rx power high alarm threshold : 0.8912 mW / -0.50 dBm
Laser rx power low alarm threshold : 0.0912 mW / -10.40 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 interfaces ge

Syntax `show interfaces ge-fpc/pic/port`
`<brief | detail | extensive | terse>`
`<media>`
`<statistics>`

Release Information Command introduced in Junos OS Release 9.0 for EX Series switches.

Description Display status information about the specified Gigabit Ethernet interface.



NOTE: You must have a transceiver plugged into an SFP or SFP+ port before information about the interface can be displayed.

Options `ge-fpc/pic/port`—Display standard information about the specified Gigabit Ethernet interface.

`brief | detail | extensive | terse`—(Optional) Display the specified level of output.

`media`—(Optional) Display media-specific information about network interfaces.

`statistics`—(Optional) Display static interface statistics.

Required Privilege Level view

- Related Documentation**
- [Monitoring Interface Status and Traffic on page 113](#)
 - [Troubleshooting Network Interfaces on EX3200 Switches on page 121](#)
 - [Troubleshooting Network Interfaces on EX4200 Switches on page 122](#)
 - [Troubleshooting an Aggregated Ethernet Interface on page 123](#)
 - [Junos OS Ethernet Interfaces Configuration Guide](#)

List of Sample Output [show interfaces ge-0/0/0 on page 235](#)
[show interfaces ge-0/0/0 brief on page 235](#)
[show interfaces ge-0/0/0 detail on page 236](#)
[show interfaces ge-0/0/4 extensive on page 237](#)

Output Fields [Table 35 on page 228](#) lists the output fields for the **show interfaces ge-** command. Output fields are listed in the approximate order in which they appear.

Table 35: show interfaces ge- Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels

Table 35: show interfaces ge- Output Fields (*continued*)

Field Name	Field Description	Level of Output
Enabled	State of the interface: Enabled or Disabled .	All levels
Interface index	Index number of the physical interface, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP index number for the physical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Description	Optional user-specified description.	brief detail extensive
Link-level type	Encapsulation being used on the physical interface.	All levels
MTU	Maximum transmission unit size on the physical interface. Default is 1514.	All levels
Speed	Speed of the interface: Auto if autonegotiation of speed is enabled; speed in megabits per second if the interface speed is explicitly configured.	All levels
Duplex	Link mode of interface: Auto if autonegotiation of link mode is enabled; Full-Duplex or Half-Duplex if the link mode is explicitly configured.	All levels
Loopback	Loopback status: Enabled or Disabled . If loopback is enabled, type of loopback: Local or Remote .	All levels
Source filtering	Source filtering status: Enabled or Disabled .	All levels
Flow control	Flow control status: Enabled or Disabled .	All levels
Auto-negotiation	Autonegotiation status: Enabled or Disabled .	All levels
Remote-fault	Remote fault status: <ul style="list-style-type: none"> • Online—Autonegotiation is manually configured as online. • Offline—Autonegotiation is manually configured as offline. 	All levels
Device flags	Information about the physical device.	All levels
Interface flags	Information about the interface.	All levels
Link flags	Information about the link.	All levels
CoS queues	Number of CoS queues configured.	detail extensive none
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	detail extensive
Current address	Configured MAC address.	detail extensive none
Hardware address	MAC address of the hardware.	detail extensive none

Table 35: show interfaces ge- Output Fields (*continued*)

Field Name	Field Description	Level of Output
Last flapped	Date, time, and how long ago the interface went from down to up. The format is Last flapped: year-month-day hour:minute:second timezone (hour:minute:second ago) . For example, Last flapped: 2008-01-16 10:52:40 UTC (3d 22:58 ago) .	detail extensive none
Statistics last cleared	Time when the statistics for the interface were last set to zero.	detail extensive
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface • Output packets—Number of packets transmitted on the interface. <p>NOTE: The bandwidth bps counter is not enabled on the switch.</p>	detail extensive
Input errors	<p>Input errors on the interface. The following paragraphs explain the counters whose meaning might 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 number 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 the 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 on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • Resource errors—Sum of transmit drops. 	extensive

Table 35: show interfaces ge- Output Fields (*continued*)

Field Name	Field Description	Level of Output
Output errors	<p>Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • Carrier transitions—Number of times the interface has gone from down to up. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning. • Errors—Sum of the outgoing frame aborts and FCS errors. • Drops—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Collisions—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number should always remain 0. If it is nonzero, there is a software bug. • Aged packets—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware. • FIFO errors—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • HS link CRC errors—Number of errors on the high-speed links between the ASICs responsible for handling the switch interfaces. • MTU errors—Number of packets whose size exceeded the MTU of the interface. • Resource errors—Sum of transmit drops. 	extensive
Egress queues	Total number of egress queues supported on the specified interface.	detail extensive
Queue counters (Egress)	<p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> • Queued packets—Number of queued packets. • Transmitted packets—Number of transmitted packets. • Dropped packets—Number of packets dropped by the ASIC's RED mechanism. 	detail extensive
Active alarms and Active defects	<p>Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. Based on the switch configuration, an alarm can ring the red or yellow alarm bell on the switch or turn on the red or yellow alarm LED on the front of the switch. These fields can contain the value None or Link.</p> <ul style="list-style-type: none"> • None—There are no active defects or alarms. • Link—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning. 	detail extensive none

Table 35: show interfaces ge- Output Fields (*continued*)

Field Name	Field Description	Level of Output
MAC statistics	<p>Receive and Transmit statistics reported by the PIC's MAC subsystem.</p> <ul style="list-style-type: none"> • Total octets and total packets—Total number of octets and packets. For Gigabit Ethernet IQ PICs, the received octets count varies by interface type. • Unicast packets, Broadcast packets, and Multicast packets—Number of unicast, broadcast, and multicast packets. • CRC/Align errors—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error). • FIFO error—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • MAC control frames—Number of MAC control frames. • MAC pause frames—Number of MAC control frames with pause operational code. • Oversized frames—Number of frames that exceed 1518 octets. • Jabber frames—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms. • Fragment frames—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted. • Code violations—Number of times an event caused the PHY to indicate "Data reception error" or "invalid data symbol error." 	extensive
Filter Statistics	Receive and Transmit statistics reported by the PIC's MAC address filter subsystem.	extensive

Table 35: show interfaces ge- Output Fields (*continued*)

Field Name	Field Description	Level of Output
Autonegotiation information	<p>Information about link autonegotiation:</p> <ul style="list-style-type: none"> • Negotiation status: <ul style="list-style-type: none"> • Complete—The autonegotiation process between the local and remote Ethernet interfaces was successful. • Incomplete—Remote Ethernet interface has the speed or link mode configured or does not perform autonegotiation. • No autonegotiation—Local Ethernet interface has autonegotiation disabled and the link mode and speed are manually configured. • Link partner—Information from the link partner: <ul style="list-style-type: none"> • Link mode—Depending on the capability of the attached Ethernet device, either Full-duplex or Half-duplex. If the link mode of the remote device cannot be determined, value is Unknown. • Flow control—Types of flow control supported by the remote Ethernet device. For Gigabit Ethernet interfaces, types are Symmetric (link partner supports PAUSE on receive and transmit), Asymmetric (link partner supports PAUSE on transmit), and Symmetric/Asymmetric (link partner supports PAUSE on both receive and transmit or PAUSE only on receive). • Remote fault—Remote fault information from the link partner—Failure indicates a receive link error. OK indicates that the link partner is receiving. Negotiation error indicates a negotiation error. Offline indicates that the link partner is going offline. • Link partner speed—Speed of the link partner. • Local resolution—Resolution of the autonegotiation process on the local interface: <ul style="list-style-type: none"> • Flow control—Type of flow control that is used by local interface. For Gigabit Ethernet interfaces, types are Symmetric (link partner supports PAUSE on receive and transmit), Asymmetric (link partner supports PAUSE on transmit), and Symmetric/Asymmetric (link partner supports PAUSE on both receive and transmit or PAUSE only on receive). • Link mode—Link mode of local interface: either Full-duplex or Half-duplex. Displayed when Negotiation status is Incomplete. • Local link speed—Speed of the local interface. Displayed when Negotiation status is Incomplete. • Remote fault—Remote fault information. Link OK (no error detected on receive), Offline (local interface is offline), and Link Failure (link error detected on receive). 	extensive
Packet Forwarding Engine configuration	<p>Information about the configuration of the Packet Forwarding Engine:</p> <ul style="list-style-type: none"> • Destination slot—FPC slot number: <ul style="list-style-type: none"> • On standalone switches with built-in interfaces, the slot number refers to the switch itself and is always 0. • On Virtual Chassis composed of switches with built-in interfaces, the slot number refers to the member ID of the switch. • On switches with line cards or on Virtual Chassis composed of switches with line cards, the slot number refers to the line card slot number on the switch or Virtual Chassis. 	extensive

Logical Interface

Table 35: show interfaces ge- Output Fields (*continued*)

Field Name	Field Description	Level of Output
Logical interface	Name of the logical interface.	All levels
Index	Index number of the logical interface, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP interface index number for the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Flags	Information about the logical interface.	All levels
Encapsulation	Encapsulation on the logical interface.	All levels
Protocol	Protocol family.	detail extensive none
Traffic statistics	<p>Number and rate of bytes and packets received (input) and transmitted (output) on the specified interface.</p> <p>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.</p>	detail extensive
IPv6 transit statistics	If IPv6 statistics tracking is enabled, number of IPv6 bytes and packets received and transmitted on the logical interface.	extensive
Local statistics	Number and rate of bytes and packets destined to and from the switch.	extensive
Transit statistics	Number and rate of bytes and packets transiting the switch.	extensive
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Route Table	Route table in which the logical interface address is located. For example, 0 refers to the routing table inet.0 .	detail extensive none
Input Filters	Names of any input filters applied to this interface.	detail extensive
Output Filters	Names of any output filters applied to this interface.	detail extensive
Flags	<p>Information about protocol family flags.</p> <p>If unicast reverse-path forwarding (RPF) is explicitly configured on the specified interface, the uRPF flag is displayed. If unicast RPF was configured on a different interface (and therefore is enabled on all switch interfaces) but was not explicitly configured on the specified interface, the uRPF flag is not displayed even though unicast RPF is enabled.</p>	detail extensive
<i>protocol-family</i>	Protocol family configured on the logical interface. If the protocol is inet , the IP address of the interface is also displayed.	brief
Flags	Information about the address flags.	detail extensive none

Table 35: show interfaces ge- Output Fields (*continued*)

Field Name	Field Description	Level of Output
Destination	IP address of the remote side of the connection.	detail extensive none
Local	IP address of the logical interface.	detail extensive none
Broadcast	Broadcast address of the logical interlace.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive

Sample Output

```

show interfaces user@switch> show interfaces ge-0/0/0
ge-0/0/0 Physical interface: ge-0/0/0, Enabled, Physical link is Down
          Interface index: 129, SNMP ifIndex: 21
          Link-level type: Ethernet, MTU: 1514, Speed: Unspecified, Loopback: Disabled,
          Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled
          Remote fault: Online
          Device flags   : Present Running Down
          Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
          CoS queues     : 8 supported, 8 maximum usable queues
          Hold-times     : Up 0 ms, Down 0 ms
          Current address: 00:19:e2:50:3f:41, Hardware address: 00:19:e2:50:3f:41
          Last flapped   : 2008-01-16 11:40:53 UTC (4d 02:30 ago)
          Input rate     : 0 bps (0 pps)
          Output rate    : 0 bps (0 pps)
          Ingress rate at Packet Forwarding Engine : 0 bps (0 pps)
          Ingress drop rate at Packet Forwarding Engine : 0 bps (0 pps)
          Active alarms  : None
          Active defects : None

          Logical interface ge-0/0/0.0 (Index 65) (SNMP ifIndex 22)
          Flags: SNMP-Traps
          Encapsulation: ENET2
          Input packets : 0
          Output packets: 0
          Protocol eth-switch
          Flags: None

show interfaces user@switch> show interfaces ge-0/0/0 brief
ge-0/0/0 brief Physical interface: ge-0/0/0, Enabled, Physical link is Down
               Description: voice priority and tcp and icmp traffic rate-limiting filter at ingress port
               Link-level type: Ethernet, MTU: 1514, Speed: Unspecified, Loopback: Disabled,
               Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
               Remote fault: Online
               Device flags   : Present Running Down
               Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
               Link flags     : None

               Logical interface ge-0/0/0.0
               Flags: Device-Down SNMP-Traps Encapsulation: ENET2
               eth-switch

```

```

show interfaces user@switch> show interfaces ge-0/0/0 detail
ge-0/0/0 detail Physical interface: ge-0/0/0, Enabled, Physical link is Up
                  Interface index: 193, SNMP ifIndex: 206, Generation: 196
                  Link-level type: Ethernet, MTU: 1514, Speed: Auto, Duplex: Auto,
                  BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
                  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
                  Remote fault: Online
                  Device flags   : Present Running
                  Interface flags: 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: 00:1f:12:30:ff:40, Hardware address: 00:1f:12:30:ff:40
                  Last flapped   : 2009-05-05 06:03:05 UTC (00:22:13 ago)
                  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
                  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 : None
                  Active defects : None

                  Logical interface ge-0/0/0.0 (Index 65) (SNMP ifIndex 235) (Generation 130)
                  Flags: SNMP-Traps Encapsulation: ENET2
                  Bandwidth: 0
                  Traffic 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
                  Protocol eth-switch, Generation: 146, Route table: 0
                  Flags: Is-Primary
                  Input Filters: f1,
                  Output Filters: f2,,,,

```

```

show interfaces user@switch> show interfaces ge-0/0/4 extensive
ge-0/0/4 extensive Physical interface: ge-0/0/4, Enabled, Physical link is Up
                    Interface index: 165, SNMP ifIndex: 152, Generation: 168
                    Link-level type: Ethernet, MTU: 1514, Speed: Auto, Duplex: 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: 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: 00:1f:12:33:65:44, Hardware address: 00:1f:12:33:65:44
                    Last flapped   : 2008-09-17 11:02:25 UTC (16:32:54 ago)
                    Statistics last cleared: Never
                    Traffic statistics:
                    Input bytes   :           0           0 bps
                    Output bytes  :      2989761      984 bps
                    Input packets :           0           0 pps
                    Output packets:      24307       1 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: 1, 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              24307            0

                    Active alarms  : None
                    Active defects : None
                    MAC statistics:
                    Total octets      Receive      Transmit
                    Total packets     0            2989761
                    Unicast packets   0            24307
                    Broadcast packets 0             0
                    Multicast packets 0            24307
                    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
                    Code violations    0
                    Autonegotiation information:
                    Negotiation status: Complete

```

```

Link partner:
  Link mode: Full-duplex, Flow control: None, Remote fault: OK,
  Link partner Speed: 1000 Mbps
Local resolution:
  Flow control: None, Remote fault: Link OK
Packet Forwarding Engine configuration:
  Destination slot: 0
  Direction : Output
  CoS transmit queue          Bandwidth          Buffer Priority
Limit
  0 best-effort              95          950000000    95          usec      low
none
  7 network-control          5           500000000    5           NA        low
none

```

```

Logical interface ge-0/0/4.0 (Index 82) (SNMP ifIndex 184) (Generation 147)
Flags: SNMP-Traps Encapsulation: ENET2

```

Traffic statistics:

```

Input bytes : 0
Output bytes : 4107883
Input packets: 0
Output packets: 24307

```

IPv6 transit statistics:

```

Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0

```

Local statistics:

```

Input bytes : 0
Output bytes : 4107883
Input packets: 0
Output packets: 24307

```

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 eth-switch, Generation: 159, Route table: 0

```

```

Flags: None

```

```

Input Filters: f2,

```

```

Output Filters: f1,,,,

```

show interfaces me0

Syntax	<pre>show interfaces me0 <brief detail extensive terse> <descriptions> <media> <routing-instance> <statistics></pre>
Release Information	Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display status information about the management Ethernet interface.
Options	<p>none—Display standard information about the management Ethernet interface.</p> <p>brief detail extensive terse—(Optional) Display the specified level of output.</p> <p>descriptions—(Optional) Display interface description strings.</p> <p>media—(Optional) Display media-specific information about network interfaces.</p> <p>routing-instance—(Optional) Display the name of the routing instance.</p> <p>statistics—(Optional) Display static interface statistics.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring a Firewall Filter on a Management Interface on an EX Series Switch • Configuring Firewall Filters (CLI Procedure)
List of Sample Output	show interfaces me0 on page 243 show interfaces me0 brief on page 243 show interfaces me0 detail on page 243 show interfaces me0 extensive on page 244
Output Fields	Table 36 on page 239 lists the output fields for the show interfaces me0 command. Output fields are listed in the approximate order in which they appear.

Table 36: show interfaces me0 Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface: Enabled or Disabled .	All levels
Interface index	Index number of the physical interface, which reflects its initialization sequence.	detail extensive none

Table 36: show interfaces me0 Output Fields (*continued*)

Field Name	Field Description	Level of Output
SNMP ifIndex	SNMP index number for the physical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Description	Optional user-specified description.	brief detail extensive
Type	Information about the type of functional interface.	All levels
Link-level type	Encapsulation being used on the physical interface.	All levels
MTU	Maximum transmission unit size on the physical interface. The default is 1514.	All levels
Clocking	Interface that acts as a clock source. This field is not supported on EX Series switches and the default value is always Unspecified .	detail extensive
Speed	Speed at which the interface is running.	All levels
Device flags	Information about the physical device.	All levels
Interface flags	Information about the interface.	All levels
Link type	Information about whether the link is duplex and whether the negotiation is manual or automatic.	detail extensive none
Physical info	Information about the device dependent physical interface selector. This field is applied only when a clocking option is specified. This field is not supported on EX Series switches and the default value is always Unspecified .	detail extensive
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	detail extensive
Current address	Configured MAC address.	detail extensive none
Hardware address	MAC address of the hardware.	detail extensive none
Alternate link address	Information about alternate hardware address.	detail extensive
Last flapped	Date, time, and how long ago the interface went from down to up. The format is Last flapped: year-month-day hour:minute:second timezone (weeksw:daysdhour:minute:second ago) . For example, Last flapped: 2008-01-16 10:52:40 UTC (3w:3d 22:58 ago) .	detail extensive none
Statistics last cleared	Time when the statistics for the interface was last set to zero. The format is Last flapped: year-month-day hour:minute:second timezone (weeksw:daysdhour:minute:second ago) . For example, Last flapped: 2008-01-16 10:52:40 UTC (3w:3d 22:58 ago) .	detail extensive

Table 36: show interfaces me0 Output Fields (*continued*)

Field Name	Field Description	Level of Output
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <p>Following are fields in Traffic statistics:</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. 	detail extensive
IPv6 transit statistics	<p>Number and rate of bytes and IPv6 packets received and transmitted on the physical interface.</p> <p>Following are fields in IPv6 transit statistics:</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. 	detail extensive
Input errors	<p>Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • Errors—Sum of the incoming frame aborts and frame checksum (FCS) errors. • Drops—Number of packets dropped by the input queue of the I/O Manager ASIC. • Framing errors—Number of packets received with an invalid FCS. • Runts—Number of frames received that are smaller than the runt threshold. • Giants—Number of packets that exceed the size for the medium. For example, if the medium is Ethernet, the Giant field shows the count of packets with size greater than 1518 bytes. • 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 the Junos OS does not handle. • Resource errors—Sum of transmit drops. 	extensive
Output errors	<p>Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • Carrier transitions—Number of times the interface has gone from down to up. This number does not normally increment quickly. It increases only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increment quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning. • Errors—Sum of the outgoing frame aborts and FCS errors. • Drops—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • MTU errors—Number of packets whose size exceeded the MTU of the interface. • Resource errors—Sum of transmit drops. 	extensive

Table 36: show interfaces me0 Output Fields (*continued*)

Field Name	Field Description	Level of Output
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Index number of the logical interface, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP interface index number for the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Flags	Information about the logical interface.	All levels
Encapsulation	Encapsulation on the logical interface.	All levels
Traffic statistics	Number and rate of bytes and packets received (input) and transmitted (output) on the specified interface.	detail extensive
IPv6 transit statistics	If IPv6 statistics tracking is enabled, number of IPv6 bytes and packets received and transmitted on the logical interface.	detail extensive
Local statistics	Number and rate of bytes and packets destined to and exiting from the switch.	extensive
Protocol	Protocol family.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Route Table	Routing table in which the logical interface address is located. For example, 0 refers to the routing table inet.0 .	detail extensive
Flags	Information about protocol family flags.	detail extensive
Input Filter	Ingress filter name.	extensive
Output Filter	Egress filter name.	extensive
Addresses	Information about the management interface addresses.	detail extensive none
Flags	Information about the address flags.	detail extensive none
Destination	IP address of the remote side of the connection.	detail extensive none
Local	IP address of the logical interface.	detail extensive none
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 me0 user@switch> show interfaces me0
Physical interface: me0, Enabled, Physical link is Up
  Interface index: 1, SNMP ifIndex: 33
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Current address: 00:1f:12:35:3c:bf, Hardware address: 00:1f:12:35:3c:bf
  Last flapped   : 2010-07-31 23:45:50 PDT (5d 00:32 ago)
    Input packets : 1661830
    Output packets: 3200

Logical interface me0.0 (Index 3) (SNMP ifIndex 34)
  Flags: SNMP-Traps Encapsulation: ENET2
  Input packets : 1661830
  Output packets: 3200
  Protocol inet
    Flags: Is-Primary
    Addresses, Flags: Is-Preferred Is-Primary
      Destination: 10.204.32/20, Local: 10.204.33.103,
      Broadcast: 10.204.47.255
  Protocol inet6
    Flags: Is-Primary
    Addresses, Flags: Is-Preferred
      Destination: fe80::21f:12ff:fe35:3cbf

show interfaces me0 user@switch> show interfaces me0 brief
brief Physical interface: me0, Enabled, Physical link is Up
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
  Speed: 1000mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps

Logical interface me0.0
  Flags: SNMP-Traps Encapsulation: ENET2
  inet 10.204.33.103/20
  inet6 fe80::21f:12ff:fe35:3cbf/64

show interfaces me0 user@switch> show interfaces me0 detail
detail Physical interface: me0, Enabled, Physical link is Up
  Interface index: 1, SNMP ifIndex: 33, Generation: 1
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
  Speed: 1000mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Physical info   : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:1f:12:35:3c:bf, Hardware address: 00:1f:12:35:3c:bf
  Alternate link address: Unspecified
  Last flapped   : 2010-07-31 23:45:50 PDT (5d 00:37 ago)
  Statistics last cleared: Never
  Traffic statistics:
    Input bytes : 366663167
    Output bytes : 498590
    Input packets: 1664031
    Output packets: 3259

```

IPv6 transit statistics:

```

Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0

```

Logical interface me0.0 (Index 3) (SNMP ifIndex 34) (Generation 1)

Flags: SNMP-Traps Encapsulation: ENET2

Traffic statistics:

```

Input bytes : 366665637
Output bytes : 500569
Input packets: 1664048
Output packets: 3275

```

IPv6 transit statistics:

```

Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0

```

Local statistics:

```

Input bytes : 366665637
Output bytes : 500569
Input packets: 1664048
Output packets: 3275

```

Protocol inet, Generation: 1, Route table: 0

Flags: Is-Primary

Addresses, Flags: Is-Preferred Is-Primary

Destination: 10.204.32/20, Local: 10.204.33.103, Broadcast: 10.204.47.255,

Generation: 1

Protocol inet6, Generation: 2, Route table: 0

Flags: Is-Primary

Addresses, Flags: Is-Preferred

Destination: fe80::/64, Local: fe80::21f:12ff:fe35:3cbf

Generation: 2

show interfaces me0
extensive

user@switch> show interfaces me0 extensive

Physical interface: me0, Enabled, Physical link is Up

Interface index: 1, SNMP ifIndex: 33, Generation: 1

Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified, Speed: 100mbps

Device flags : Present Running

Interface flags: SNMP-Traps

Link type : Full-Duplex

Physical info : Unspecified

Hold-times : Up 0 ms, Down 0 ms

Current address: 00:1f:12:38:58:bf, Hardware address: 00:1f:12:38:58:bf

Alternate link address: Unspecified

Last flapped : 2010-08-15 06:27:33 UTC (03:06:22 ago)

Statistics last cleared: Never

Traffic statistics:

```

Input bytes : 82310392
Output bytes : 1966952
Input packets: 110453
Output packets: 17747

```

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, Giants: 0,
Policed discards: 0, Resource errors: 0

```

Output errors:

Carrier transitions: 1, Errors: 0, Drops: 0, MTU errors: 0,
Resource errors: 0

Logical interface me0.0 (Index 3) (SNMP ifIndex 34) (Generation 1)

Flags: SNMP-Traps Encapsulation: ENET2

Traffic statistics:

Input bytes :	82310392
Output bytes :	1966952
Input packets:	110453
Output packets:	17747

Local statistics:

Input bytes :	82310392
Output bytes :	1966952
Input packets:	110453
Output packets:	17747

Protocol inet, Generation: 1, Route table: 0

Flags: Is-Primary

Input Filters: mgmt_filter,

Addresses, Flags: Is-Default Is-Preferred Is-Primary

Destination: 10.204.96/20, Local: 10.204.96.234,

Broadcast: 10.204.111.255, Generation: 1

show interfaces queue

Syntax	show interfaces queue <both-ingress-egress> <egress> <forwarding-class <i>forwarding-class</i> > <ingress> < <i>interface-name</i> >
Release Information	Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display class-of-service (CoS) queue information for physical interfaces.
Options	<p>none—Show detailed CoS queue statistics for all physical interfaces.</p> <p>both-ingress-egress—(Optional) Show both ingress and egress queue statistics. (Ingress statistics are not available for all interfaces.)</p> <p>egress—(Optional) Show egress queue statistics only.</p> <p>forwarding-class <i>forwarding-class</i>—(Optional) Show queue statistics only for the specified forwarding class.</p> <p>ingress—(Optional) Show ingress queue statistics only. (Ingress statistics are not available for all interfaces.)</p> <p><i>interface-name</i>—(Optional) Show queue statistics for the specified interface.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Monitoring Interface Status and Traffic on page 113 • Monitoring Interfaces That Have CoS Components • Defining CoS Schedulers (CLI Procedure) • Configuring CoS Traffic Classification for Ingress Queuing on Oversubscribed Ports on EX8200 Line Cards (CLI Procedure)
List of Sample Output	<p>show interfaces queue ge-0/0/0 (EX2200 Switch) on page 248</p> <p>show interfaces queue xe-6/0/39 (Line Card with Oversubscribed Ports in an EX8200 Switch) on page 249</p>
Output Fields	<p>Table 37 on page 246 lists the output fields for the show interfaces queue command. Output fields are listed in the approximate order in which they appear.</p>

Table 37: show interfaces queue Output Fields

Field Name	Field Description
Physical Interface and Forwarding Class Information	
Physical interface	Name of the physical interface.

Table 37: show interfaces queue Output Fields (*continued*)

Field Name	Field Description
Enabled	<p>State of the interface. Possible values are:</p> <ul style="list-style-type: none"> • Administratively down, Physical link is Down—The interface is turned off, and the physical link is inoperable. • Administratively down, Physical link is Up—The interface is turned off, but the physical link is operational and can pass packets when it is enabled. • Enabled, Physical link is Down—The interface is turned on, but the physical link is inoperable and cannot pass packets. • Enabled, Physical link is Up—The interface is turned on, and the physical link is operational and can pass packets.
Interface index	Index number of the physical interface, which reflects its initialization sequence.
SNMP ifIndex	SNMP index number for the physical interface.
Description	User-configured interface description.
Forwarding classes	Number of forwarding classes supported and in use for the interface.
Ingress Queues Information (not shown for all interfaces)	
Ingress queues	Number of input queues supported and in use on the specified interface. For an interface on a line card with oversubscribed ports, the ingress queue handles low priority traffic on the interface.
Transmitted	<p>Transmission statistics for the queue:</p> <ul style="list-style-type: none"> • Packets—Number of packets transmitted by this queue. • Bytes—Number of bytes transmitted by this queue. • Tail-dropped packets—Number of packets dropped because the queue buffers were full.
PFE chassis queues	For an interface on a line card with oversubscribed ports, the number of Packet Forwarding Engine chassis queues supported and in use for the port group to which the interface belongs. The Packet Forwarding Engine chassis queue for a port group handles high priority traffic from all the interfaces in the port group.
Egress Queues Information	
Egress queues	Number of output queues supported and in use on the specified interface.
Queue	CoS queue number.
Queued	This counter is not supported on EX Series switches.

Table 37: show interfaces queue Output Fields (*continued*)

Field Name	Field Description
Transmitted	<p>Number of packets and bytes transmitted by this queue. Information on transmitted packets and bytes can include:</p> <ul style="list-style-type: none"> • Packets—Number of packets transmitted. • Bytes—Number of bytes transmitted. • Tail-dropped packets—Number of arriving packets dropped because output queue buffers were full. • RED-dropped packets—Number of packets dropped because of random early detection (RED). <ul style="list-style-type: none"> • Low—Number of low loss priority packets dropped because of RED. • High—Number of high loss priority packets dropped because of RED. • RED-dropped bytes—Number of bytes dropped because of random early detection (RED). <ul style="list-style-type: none"> • Low—Number of low loss priority bytes dropped because of RED. • High—Number of high loss priority bytes dropped because of RED.
Packet Forwarding Engine Chassis Queues	<p>For an interface on a line card with oversubscribed ports, the number of Packet Forwarding Engine chassis queues supported and in use for the port group to which the interface belongs. The queue statistics reflect the traffic flowing on all the interfaces in the port group.</p>

Sample Output

```

show interfaces queue ge-0/0/0 (EX2200 Switch)
user@switch> show interfaces queue ge-0/0/0
Physical interface: ge-0/0/0, Enabled, Physical link is Down
  Interface index: 130, SNMP ifIndex: 501
  Forwarding classes: 16 supported, 4 in use
  Egress queues: 8 supported, 4 in use
  Queue: 0, Forwarding classes: best-effort
    Queued:
    Transmitted:
      Packets      : 0
      Bytes       : 0
      Tail-dropped packets : 0
  Queue: 1, Forwarding classes: assured-forwarding
    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: 7, Forwarding classes: network-control
    Queued:
    Transmitted:
      Packets      : 0

```

```

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:
        Packets : 46714
        Bytes : 6901326

```

```
Tail-dropped packets : 0

Packet Forwarding Engine Chassis Queues:
Queues: 8 supported, 7 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Transmitted:
      Packets : 739338141426
      Bytes : 94635282101928
      Tail-dropped packets : 0
      RED-dropped packets : 5606426444
        Low : 5606426444
        High : 0
      RED-dropped bytes : 683262846464
        Low : 683262846464
        High : 0
  Queue: 1, Forwarding classes: assured-forwarding
    Queued:
      Transmitted:
        Packets : 0
        Bytes : 0
        Tail-dropped packets : 0
        RED-dropped packets : 0
          Low : 0
          High : 0
        RED-dropped bytes : 0
          Low : 0
          High : 0
  Queue: 2, Forwarding classes: mcast-be
    Queued:
      Transmitted:
        Packets : 0
        Bytes : 0
        Tail-dropped packets : 0
        RED-dropped packets : 0
          Low : 0
          High : 0
        RED-dropped bytes : 0
          Low : 0
          High : 0
  Queue: 4, Forwarding classes: mcast-ef
    Queued:
      Transmitted:
        Packets : 0
        Bytes : 0
        Tail-dropped packets : 0
        RED-dropped packets : 0
          Low : 0
          High : 0
        RED-dropped bytes : 0
          Low : 0
          High : 0
  Queue: 5, Forwarding classes: expedited-forwarding
    Queued:
      Transmitted:
        Packets : 0
        Bytes : 0
        Tail-dropped packets : 0
        RED-dropped packets : 0
          Low : 0
          High : 0
```



```
RED-dropped bytes : 0
  Low : 0
  High : 0
Queue: 6, Forwarding classes: mcast-af
Queued:
Transmitted:
  Packets : 0
  Bytes : 0
Tail-dropped packets : 0
RED-dropped packets : 0
  Low : 0
  High : 0
RED-dropped bytes : 0
  Low : 0
  High : 0
Queue: 7, Forwarding classes: network-control
Queued:
Transmitted:
  Packets : 97990
  Bytes : 14987506
Tail-dropped packets : 0
RED-dropped packets : 0
  Low : 0
  High : 0
RED-dropped bytes : 0
  Low : 0
  High : 0
```

show interfaces vlan

Syntax	<code>show interfaces (vlan vlan.vlan-id)</code> <code><brief detail extensive terse></code> <code><descriptions></code> <code><media></code> <code><routing-instance (all instance-name)></code> <code><snmp-index snmp-index></code> <code><statistics></code>
Release Information	Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display status information about routed VLAN interfaces (RVIs).
Options	<code>vlan vlan.vlan-id</code> —Display status information for the specified RVI. <code>brief detail extensive terse</code> —(Optional) Display the specified level of output. <code>descriptions</code> —(Optional) Display interface description strings. <code>media</code> —(Optional) Display media-specific information about network interfaces. <code>routing-instance (all instance-name)</code> —(Optional) Associate this RVI with the named routing instance. <code>snmp-index snmp-index</code> —(Optional) Display information for the specified SNMP index of the interface. <code>statistics</code> —(Optional) Display static interface statistics.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• show ethernet-switching table• show vlans• Monitoring Interface Status and Traffic on page 113• Troubleshooting Network Interfaces on EX3200 Switches on page 121• Troubleshooting Network Interfaces on EX4200 Switches on page 122• Verifying Routed VLAN Interface Status and Statistics
List of Sample Output	show interfaces vlan on page 260 show interfaces vlan terse on page 260 show interfaces vlan extensive on page 261 show interfaces vlan detail on page 262
Output Fields	Table 38 on page 253 lists the output fields for the show interfaces vlan command. Output fields are listed in the approximate order in which they appear. The level of output none means the basic command with no optional options—that is, either just show interfaces vlan or show interfaces vlan.vlan-id .

Table 38: show interfaces vlan Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface, which is always vlan .	All levels
Enabled	State of the interface: Enabled or Disabled , followed by the statement Physical link is <Up/Down>	All levels
Interface index	Index number of the physical interface, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP index number for the physical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Type	Because this is routed VLAN interface information, this entry is always VLAN .	detail extensive none
Link-level type	Encapsulation (added control information) being used on the physical interface. Because this is routed VLAN interface information, this entry is always VLAN .	All levels
MTU	Maximum transmission unit (MTU) size on the physical interface. The default MTU size depends on the switch platform. Changing either the media MTU or protocol MTU causes an interface to be deleted and added again.	All levels
Clocking	Value is always Unspecified —not applicable on switches.	detail extensive
Speed	Speed of the interface, either Auto if autonegotiation of speed is enabled or a number representing the configured speed in megabits per second.	detail extensive none

Table 38: show interfaces vlan Output Fields (*continued*)

Field Name	Field Description	Level of Output
Device flags	<p>Information about the physical device such as:</p> <p>Dest-route-down—The routing process detected that the link was not operational and changed the interface routes to nonforwarding status.</p> <p>Down—Device has been administratively disabled.</p> <p>Hear-Own-Xmit—Device receives its own transmissions.</p> <p>Is-Default—This address is the default address of the switch. The default address is used as the source address by SNMP, ping, traceroute, and other network utilities.</p> <p>Is-Preferred—This address is the default local address for packets originating from the local switch and sent to destinations on the subnet.</p> <p>Is-Primary—This address is the default local address for broadcast and multicast packets originated locally and sent out the interface.</p> <p>Link-Layer-Down—The link-layer protocol has failed to connect with the remote endpoint.</p> <p>Loopback—Switch is in physical loopback.</p> <p>Loop-Detected—The link layer has received frames that it sent, thereby detecting a physical loopback.</p> <p>No-Carrier—On media that support carrier recognition, no carrier is currently detected.</p> <p>No-Multicast—Device does not support multicast traffic.</p> <p>Preferred—This address is a candidate to become the preferred address.</p> <p>Present—Device is physically present and recognized.</p> <p>Promiscuous—Device is in promiscuous mode and recognizes frames addressed to all physical addresses on the media.</p> <p>Primary—This address is a candidate to become the primary address.</p> <p>Quench—Transmission on the device is quenched, because the output buffer is overflowing.</p> <p>Recv-All-Multicasts—Device is in multicast promiscuous mode and therefore provides no multicast filtering.</p> <p>Running—Device is active and enabled.</p>	<p>detail</p> <p>extensive</p> <p>none</p>
Link type	Link mode of the interface— Auto if autonegotiation is enabled, or the configured Full-Duplex or Half-Duplex .	<p>detail</p> <p>extensive</p> <p>none</p>
Link flags	Value is always None —not applicable on switches.	<p>detail</p> <p>extensive</p> <p>none</p>
Physical Info	Value is always Unspecified —not applicable on switches.	<p>detail</p> <p>extensive</p>
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	<p>detail</p> <p>extensive</p>
Current address	MAC address of the hardware.	<p>detail</p> <p>extensive</p> <p>none</p>
Hardware address	MAC address of the switch.	<p>detail</p> <p>extensive</p> <p>none</p>

Table 38: show interfaces vlan Output Fields (*continued*)

Field Name	Field Description	Level of Output
Alternate link address	Value is always Unspecified —not applicable on switches.	detail extensive
Last flapped	Date, time, and how long ago the interface went from down to up. The format is Last flapped: year-month-day hour:minute:second timezone (hour:minute:second ago) . For example, Last flapped: 2008-01-16 10:52:40 UTC (3d 22:58 ago) . The entry can also be Never .	detail extensive none
Statistics last cleared	Time when the statistics for the interface were last set to zero.	detail extensive none
Traffic statistics	<p>Number and rate of bytes and packets transmitted or received on the physical interface for supported switches.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface for this switch. This value reflects the information gathered by the automatic ingress counter on EX3200 switches and EX4200 switches. EX8200 switches can also be configured to collect this information with the command <code>l3-interface-ingress-counting</code>. • Output bytes—Number of bytes sent on the interface. This value reflects the information gathered by the automatic egress counter for EX8200 switches. • Input packets—Number of packets received on the interface for this switch. This value reflects the information gathered by the automatic ingress counter for EX3200 and EX4200 switches. EX8200 switches can also be configured to collect this information with the command <code>l3-interface-ingress-counting</code>. • Output packets—Number of packets sent on the interface. This value reflects the information gathered by the automatic egress counter for EX8200 switches. 	detail extensive
IPv6 transit statistics	<p>Number and rate of bytes and packets transmitted and/or received on the IPv6 interface for supported switches.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. This value reflects the information gathered by the automatic ingress counter for EX3200 and EX4200 switches. EX8200 switches can also be configured to collect this information with the command <code>l3-interface-ingress-counting</code>. • Output bytes—Number of bytes sent on the IPv6 interface. This value reflects the information gathered by the automatic egress counter for EX8200 switches. • Input packets—Number of packets received on the interface. This value reflects the information gathered by the automatic ingress counter for EX3200 and EX4200 switches. EX8200 switches can also be configured to collect this information with the command <code>l3-interface-ingress-counting</code>. • Output packets—Number of packets sent on the IPv6 interface. This value reflects the information gathered by the automatic egress counter for and EX8200 switches. 	detail extensive

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

Table 38: show interfaces vlan Output Fields (*continued*)

Field Name	Field Description	Level of Output
Output errors	<p>Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • Carrier transitions—Number of times the interface has gone from down to up. This value does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the interface is malfunctioning. • Errors—Sum of the outgoing frame aborts and FCS errors. • Drops—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Collisions—Number of Ethernet collisions. Both Gigabit Ethernet interfaces and 10 Gigabit Ethernet interfaces support only full-duplex operation, so for those two interfaces, this value should always be zero. If the value is nonzero for either Gigabit Ethernet or 10 Gigabit Ethernet, there is a software bug. • Aged packets—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware. • FIFO errors—Number of FIFO errors in the send direction as reported by the ASIC on the interface. If this value is ever nonzero, the interface is probably malfunctioning. • HS link CRC errors—Number of errors on the high-speed links between the ASICs responsible for handling the switch interfaces. • MTU errors—Number of packets whose size exceeded the MTU of the interface. • Resource errors—Sum of transmit drops. 	extensive
Logical Interface		
vlan.vlan-id, Index, SNMP ifIndex	VLAN ID, index, and SNMP index number for the logical interface. The logical interface index values reflect the item's initialization sequence.	detail extensive none
Generation	Unique number for Juniper Networks Technical support use only.	detail extensive none
Flags	<p>Errors that have occurred on this interface, such as Link Layer Down. Other possible flags include:</p> <ul style="list-style-type: none"> • Device-down—Device has been administratively disabled. • Disabled—Interface is administratively disabled. • Down—A hardware failure has occurred. • Hardware-Down—Interface protocol initialization failed to complete successfully. • SNMP-Traps—SNMP trap notifications are enabled. • Up—Interface is enabled and operational. 	detail extensive none

Table 38: show interfaces vlan Output Fields (*continued*)

Field Name	Field Description	Level of Output
SNMP-Traps	Each configured SNMP trap has a number that appears here—0x0 is always displayed for logical interface SNMP traps.	detail extensive none
Encapsulation	Encapsulation method, which is the process of adding control information. The value is always Ethernet 2 (ENET2) for logical encapsulation.	detail extensive none
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the logical interface of supported switches. Traffic statistics represent the sum of the next two fields, Local statistics and Transit statistics. Note that these are not the values for the RVI ingress or egress counters—for that value, see Transit statistics below.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. Same value as the physical interface. • Output bytes—Number of bytes sent on the interface. Same value as the physical interface. • Input packets—Number of packets received on the interface. Same value as the physical interface. • Output packets—Number of bytes sent on the interface. Same value as the physical interface. <p>NOTE: The bandwidth bps counter is not enabled on the switches.</p>	detail extensive
Local statistics	<p>Number and rate of bytes and packets received and transmitted locally by the Routing Engine on the logical interface of supported switches. All packets for protocols and process statistics are counted here.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. Same value as for the physical interface. • Output bytes—Number of bytes sent on the interface. Same value as for the physical interface. • Input packets—Number of packets received on the interface. Same value as for the physical interface. • Output packets—Number of bytes sent on the interface. Same value as for the physical interface. 	detail extensive none
Transit statistics	<p>Number and rate of bytes and packets received and transmitted on the RVI logical interface of supported switches. Look at this value to see the RVI ingress and egress count.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. This ingress counter is automatic for EX3200 and EX4200 switches and configurable for EX8200 switches. • Output bytes—Number of bytes sent on the interface. This egress counter is automatic for EX8200. • Input packets—Number of packets received on the interface. This ingress counter is automatic for EX3200 and EX4200 switches and configurable for EX8200 switches. • Output packets—Number of packets sent on the interface. This egress counter is automatic for EX8200 switches. 	detail extensive

Table 38: show interfaces vlan Output Fields (*continued*)

Field Name	Field Description	Level of Output
IPv6 transit statistics	<p>Number and rate of IPv6 bytes and packets received and transmitted on the RVI logical interface of supported switches. Transit values are unique to the logical interface and do not appear in physical interface output. Look at the values listed below to see the RVI ingress and egress count for IPv6 traffic.</p> <ul style="list-style-type: none"> Input bytes—Number of bytes received on the interface. This ingress counter is automatic for EX3200 and EX4200 switches and configurable for EX8200 switches. Output bytes—Number of bytes sent by the interface. This egress counter is automatic for EX8200 switches. Input packets—Number of packets received on the interface. This ingress counter is automatic for EX3200 and EX4200 and configurable for EX8200 switches. Output packets—Number of packets sent by the interface. This egress counter is automatic for EX8200 switches. <p>NOTE: The bandwidth bps counter is not enabled on the switches.</p>	<p>detail extensive</p>
Protocol	Protocol used for the logical interface—this value is inet for IPv4 traffic and inet6 for IPv6 traffic.	All levels
Generation	Unique number for use by Juniper Networks technical support only.	<p>detail extensive</p>
Route table	Route table in which the logical interface address is located. For example, 0 refers to the routing table inet.0 .	<p>detail extensive none</p>
Protocol flags	Information about the protocol such as Targeted-broadcast .	<p>detail extensive none</p>
Protocol addresses and Address flags	<p>Protocol address values here can be:</p> <p>Dest-route-down—The routing process detected that the link was not operational and changed the interface routes to nonforwarding status</p> <p>Device-down—Device has been administratively disabled.</p> <p>Disabled—Interface is administratively disabled.</p> <p>Down—A hardware failure has occurred.</p> <p>Hardware-Down—Interface protocol initialization failed to complete successfully.</p> <p>Is-Default—This address is the default address of the switch. The default address is used as the source address by SNMP, ping, traceroute, and other network utilities.</p> <p>Is-Preferred—This address is the default local address for packets originating from the local switch and sent to destinations on the subnet.</p> <p>Is-Primary—This address is the default local address for broadcast and multicast packets originated locally and sent out the interface.</p> <p>Preferred—This address is a candidate to become the preferred address.</p> <p>Primary—This address is a candidate to become the primary address.</p> <p>SNMP-Traps—SNMP trap notifications are enabled.</p> <p>Up—Interface is enabled and operational.</p>	<p>detail extensive none</p>

Table 38: show interfaces vlan Output Fields (*continued*)

Field Name	Field Description	Level of Output
Address destination	Logical destination's network address.	detail extensive none
Local address	IP address of the logical interface.	detail extensive none
Broadcast address	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 vlan user@switch> show interfaces vlan
Physical interface: vlan, Enabled, Physical link is Up
Interface index: 150, SNMP ifIndex: 556
Type: VLAN, Link-level type: VLAN, MTU: 1518, Speed: 1000mbps
Device flags   : Present Running
Link type      : Full-Duplex
Link flags     : None
Current address: 00:21:59:c5:f0:40, Hardware address: 00:21:59:c5:f0:40
Last flapped   : Never
  Input packets : 0
  Output packets: 0

Logical interface vlan.0 (Index 82) (SNMP ifIndex 557)
Flags: Link-Layer-Down SNMP-Traps 0x0 Encapsulation: ENET2
Input packets : 0
Output packets: 1
Protocol inet
  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

Logical interface vlan.1 (Index 83) (SNMP ifIndex 558)
Flags: Link-Layer-Down SNMP-Traps 0x0 Encapsulation: ENET2
Input packets : 0
Output packets: 1
Protocol inet
  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

show interfaces vlan terse user@switch> show interfaces vlan terse
Interface      Admin Link Proto  Local      Remote
vlan           up      up

```

```

vlan.0          up    down inet    10.1.1.1/24
vlan.1          up    down inet    10.1.2.1/24

```

**show interfaces vlan
extensive**

```

user@switch> show interfaces vlan extensive
Physical interface: vlan, Enabled, Physical link is Up
Interface index: 150, SNMP ifIndex: 556, Generation: 153
Type: VLAN, Link-level type: VLAN, MTU: 1518, Clocking: Unspecified,
Speed: 1000mbps
Device flags   : Present Running
Link type      : Full-Duplex
Link flags     : None
Physical info  : Unspecified
Hold-times    : Up 0 ms, Down 0 ms
Current address: 00:21:59:c5:f0:40, Hardware address: 00:21:59:c5:f0:40
Alternate link address: Unspecified
Last flapped   : Never
Statistics last cleared: Never
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
Input errors:
Errors: 0, Drops: 0, Framing errors: 0, Runt: 0, Giants: 0,
Policed discards: 0, Resource errors: 0
Output errors:
Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0,
Resource errors: 0

Logical interface vlan.0 (Index 82) (SNMP ifIndex 557) (Generation 147)
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: 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

```

```

show interfaces vlan user@switch> show interfaces vlan detail
detail
Physical interface: vlan, Enabled, Physical link is Up
Interface index: 150, SNMP ifIndex: 556, Generation: 153
Type: VLAN, Link-level type: VLAN, MTU: 1518, Clocking: Unspecified,
Speed: 1000mbps
Device flags : Present Running
Link type : Full-Duplex
Link flags : None
Physical info : Unspecified
Hold-times : Up 0 ms, Down 0 ms
Current address: 00:21:59:c5:f0:40, Hardware address: 00:21:59:c5:f0:40
Alternate link address: Unspecified
Last flapped : Never
Statistics last cleared: Never
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

Logical interface vlan.0 (Index 82) (SNMP ifIndex 557) (Generation 147)
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: 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
```

show interfaces xe-


Syntax	show interfaces <i>xe-fpc/pic/port</i> <brief detail extensive terse> <media> <statistics>
Release Information	Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display status information about the specified 10-Gigabit Ethernet interface.
	<div>  <p>NOTE: You must have a transceiver plugged into an SFP+ or XFP port before information about the interface can be displayed.</p> </div>
Options	<p><i>xe-fpc/pic/port</i>—Display standard information about the specified 10-Gigabit Ethernet interface.</p> <p>brief detail extensive terse—(Optional) Display the specified level of output.</p> <p>media—(Optional) Display media-specific information about network interfaces. For 10-Gigabit Ethernet interfaces, using the media option does not provide you with new or additional information. The output is the same as when the media option is not used.</p> <p>statistics—(Optional) Display static interface statistics. For 10-Gigabit Ethernet interfaces, using the statistics option does not provide you with new or additional information. The output is the same as when the statistics option is not used.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Monitoring Interface Status and Traffic on page 113 • Troubleshooting Network Interfaces on EX3200 Switches on page 121 • Troubleshooting Network Interfaces on EX4200 Switches on page 122 • Troubleshooting an Aggregated Ethernet Interface on page 123 • Junos OS Ethernet Interfaces Configuration Guide
List of Sample Output	show interfaces xe-4/1/0 on page 273 show interfaces xe-0/1/0 brief on page 273 show interfaces xe-4/1/0 detail on page 273 show interfaces xe-6/0/39 extensive on page 274
Output Fields	Table 39 on page 265 lists the output fields for the show interfaces xe- command. Output fields are listed in the approximate order in which they appear.

Table 39: show interfaces xe- Output Fields

Field Name	Field Description	Level of Output
Fields for the Terse Output Level Only		
Interface	Name of the physical or logical interface.	terse
Admin	Administrative state of the interface.	terse
Link	State of the physical link.	terse
Proto	Protocol family configured on the logical interface.	terse
Local	Local IP address of the logical interface.	terse
Remote	Remote IP address of the logical interface.	terse
Fields for the Physical Interface		
Physical interface	Name of the physical interface.	brief detail extensive none
Enabled	State of the interface. Can be one of the following: <ul style="list-style-type: none"> • Administratively down, Physical link is Down—The interface is turned off, and the physical link is inoperable and cannot pass packets even when it is enabled. • Administratively down, Physical link is Up—The interface is turned off, but the physical link is operational and can pass packets when it is enabled. • Enabled, Physical link is Down—The interface is turned on, but the physical link is inoperable and cannot pass packets. • Enabled, Physical link is Up—The interface is turned on, and the physical link is operational and can pass packets. 	brief detail extensive none
Interface index	Index number of the physical interface, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP index number for the physical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Description	User-configured interface description.	brief detail extensive none

Table 39: show interfaces xe- Output Fields (*continued*)

Field Name	Field Description	Level of Output
Link-level type	Encapsulation being used on the physical interface.	brief detail extensive none
MTU	Maximum transmission unit size on the physical interface.	brief detail extensive none
Speed	Speed at which the interface is running.	brief detail extensive none
Duplex	Duplex mode of the interface.	brief detail extensive none
BPDU Error	Not supported on EX Series switches.	detail extensive none
MAC-REWRITE Error	Not supported on EX Series switches.	detail extensive none
Loopback	Loopback status: Enabled or Disabled . If loopback is enabled, type of loopback: Local or Remote .	brief detail extensive none
Source filtering	Source filtering status: Enabled or Disabled .	brief detail extensive none
Flow control	Flow control status: Enabled or Disabled .	brief detail extensive none
Device flags	Information about the physical device.	brief detail extensive none
Interface flags	Information about the interface.	brief detail extensive none

Table 39: show interfaces xe- Output Fields (*continued*)

Field Name	Field Description	Level of Output
Link flags	Information about the link.	brief detail extensive none
CoS queues	Number of CoS queues configured.	detail extensive none
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	detail extensive
Current address	Configured MAC address.	detail extensive none
Hardware address	Hardware MAC address.	detail extensive none
Last flapped	Date, time, and how long ago the interface went from down to up. The format is <i>year-month-day hour:minute:second timezone (weekswdaysd hours:minutes:seconds ago)</i> . For example, 2008-01-16 10:52:40 UTC (3d 22:58 ago).	detail extensive none
Input Rate	Input rate in bits per second (bps) and packets per second (pps).	none
Output Rate	Output rate in bps and pps.	none
Statistics last cleared	Date, time, and how long ago the statistics for the interface were cleared. The format is <i>year-month-day hour:minute:second timezone (weekswdaysd hours:minutes:seconds ago)</i> . For example, 2010-05-17 07:51:28 PDT (00:04:33 ago).	detail extensive
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface and rate in bits per second. • Output bytes—Number of bytes transmitted on the interface and rate in bits per second. • Input packets—Number of packets received on the interface and rate in packets per second. • Output packets—Number of packets transmitted on the interface and rate in packets per second. 	detail extensive

Table 39: show interfaces xe- Output Fields (*continued*)

Field Name	Field Description	Level of Output
IPv6 transit statistics	<p>If IPv6 statistics tracking is enabled, number of IPv6 bytes and packets received and transmitted on the logical interface:</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. 	detail extensive
Input errors	<p>Input errors on the interface:</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 number 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 the Junos OS does not handle. • L3 incompletes—Number of incoming packets discarded because they failed Layer 3 sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored if you configure the ignore-l3-incompletes statement. • 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 on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • Resource errors—Sum of transmit drops. 	extensive

Table 39: show interfaces xe- Output Fields (*continued*)

Field Name	Field Description	Level of Output
Output errors	<p>Output errors on the interface:</p> <ul style="list-style-type: none"> • Carrier transitions—Number of times the interface has gone from down to up. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning. • Errors—Sum of the outgoing frame aborts and FCS errors. • Drops—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Collisions—Number of Ethernet collisions. A 10-Gigabit Ethernet interface supports only full-duplex operation, so for 10-Gigabit Ethernet interfaces, this number should always remain 0. If it is nonzero, there is a software bug. • Aged packets—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware. • FIFO errors—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • HS link CRC errors—Number of errors on the high-speed links between the ASICs responsible for handling the switch interfaces. • MTU errors—Number of packets whose size exceeded the MTU of the interface. • Resource errors—Sum of transmit drops. 	extensive
Ingress queues	Number of CoS ingress queues supported on the specified interface. Displayed only for an interface on a line card with oversubscribed ports.	detail extensive
Egress queues	Number of CoS egress queues supported on the specified interface.	detail extensive
PFE Egress queues	Number of Packet Forwarding Engine egress queues shared by the interfaces in a port group. Displayed only for an interface on a line card with oversubscribed ports.	detail extensive
Queue counters	<p>Statistics for queues:</p> <ul style="list-style-type: none"> • Queued packets—Number of queued packets. This counter is not supported on EX switches and always contains 0. • Transmitted packets—Number of transmitted packets. • Dropped packets—Number of packets dropped by the ASIC's RED mechanism. 	detail extensive

Table 39: show interfaces xe- Output Fields (*continued*)

Field Name	Field Description	Level of Output
Active alarms and Active defects	<p>Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. Based on the switch configuration, an alarm can ring the red or yellow alarm bell on the switch or turn on the red or yellow alarm LED on the front of the switch. These fields can contain the value None or Link.</p> <ul style="list-style-type: none"> • None—There are no active defects or alarms. • Link—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning. 	<p>detail extensive none</p>
MAC statistics	<p>Receive and Transmit statistics reported by the PIC's MAC subsystem.</p> <ul style="list-style-type: none"> • Total octets and total packets—Total number of octets and packets. • Unicast packets, Broadcast packets, and Multicast packets—Number of unicast, broadcast, and multicast packets. • CRC/Align errors—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error). • FIFO error—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • MAC control frames—Number of MAC control frames. • MAC pause frames—Number of MAC control frames with pause operational code. • Oversized frames—Number of frames that exceed 1518 octets. • Jabber frames—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms. • Fragment frames—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. Fragment frames normally increment because both runs (which are normal occurrences caused by collisions) and noise hits are counted. • Code violations—Number of times an event caused the PHY to indicate "Data reception error" or "invalid data symbol error." 	extensive
Packet Forwarding Engine configuration	<p>Information about the configuration of the Packet Forwarding Engine:</p> <ul style="list-style-type: none"> • Destination slot—FPC slot number: <ul style="list-style-type: none"> • On standalone switches with built-in interfaces, the slot number refers to the switch itself and is always 0. • On Virtual Chassis composed of switches with built-in interfaces, the slot number refers to the member ID of the switch. • On switches with line cards or on Virtual Chassis composed of switches with line cards, the slot number refers to the line card slot number on the switch or Virtual Chassis. 	extensive

Table 39: show interfaces xe- Output Fields (*continued*)

Field Name	Field Description	Level of Output
CoS Information	<p>Scheduler information for the CoS egress queues on the physical interface:</p> <ul style="list-style-type: none"> • Direction—Queue direction, always Output. • CoS transmit queue—Queue number and its associated user-configured forwarding class name. • Bandwidth—Information about bandwidth allocated to the queue: <ul style="list-style-type: none"> • %—Bandwidth allocated to the queue as a percentage • bps—Bandwidth allocated to the queue in bps • Buffer—Information about buffer space allocated to the queue: <ul style="list-style-type: none"> • %—Buffer space allocated to the queue as a percentage. • usec—Buffer space allocated to the queue in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • Priority—Queue priority: low or high. • Limit—Displayed if rate limiting is configured for the queue. Possible values are none and exact. If exact is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If none is configured, the queue transmits beyond the configured bandwidth if bandwidth is available. 	extensive
Fields for Logical Interfaces		
Logical interface	Name of the logical interface.	brief detail extensive none
Index	Index number of the logical interface, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP interface index number for the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Description	User-configured description of the interface.	brief detail extensive none
Flags	Information about the logical interface.	brief detail extensive none
Encapsulation	Encapsulation on the logical interface.	brief detail extensive none

Table 39: show interfaces xe- Output Fields (*continued*)

Field Name	Field Description	Level of Output
Traffic statistics	Number and rate of bytes and packets received (input) and transmitted (output) on the specified interface. 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.	detail extensive
Local statistics	Number and rate of bytes and packets destined to and from the switch.	extensive
Transit statistics	Number and rate of bytes and packets transiting the switch.	extensive
Protocol	Protocol family.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Route Table	Route table in which the logical interface address is located. For example, 0 refers to the routing table inet.0 .	detail extensive none
Input Filters	Names of any input filters applied to this interface.	detail extensive
Output Filters	Names of any output filters applied to this interface.	detail extensive
Flags	Information about protocol family flags. If unicast reverse-path forwarding (RPF) is explicitly configured on the specified interface, the uRPF flag is displayed. If unicast RPF was configured on a different interface (and therefore is enabled on all switch interfaces) but was not explicitly configured on the specified interface, the uRPF flag is not displayed even though unicast RPF is enabled.	detail extensive
Addresses, Flags	Information about the address flags.	detail extensive none
<i>protocol-family</i>	Protocol family configured on the logical interface. If the protocol is inet , the IP address of the interface is also displayed.	brief
Flags	Information about the address flags.	detail extensive none
Destination	IP address of the remote side of the connection.	detail extensive none

Table 39: show interfaces xe- Output Fields (*continued*)

Field Name	Field Description	Level of Output
Local	IP address of the logical interface.	detail extensive none
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 user@switch show interfaces xe-4/1/0
xe-4/1/0 Physical interface: xe-4/1/0, Enabled, Physical link is Up
          Interface index: 387, SNMP ifIndex: 369
          Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex,
          BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
          Source filtering: Disabled, Flow control: Enabled
          Device flags   : Present Running
          Interface flags: SNMP-Traps Internal: 0x0
          Link flags     : None
          CoS queues     : 8 supported, 8 maximum usable queues
          Current address: 00:23:9c:03:8e:70, Hardware address: 00:23:9c:03:8e:70
          Last flapped   : 2009-05-12 08:01:04 UTC (00:13:44 ago)
          Input rate     : 36432 bps (3 pps)
          Output rate    : 0 bps (0 pps)
          Active alarms  : None
          Active defects : None

          Logical interface xe-4/1/0.0 (Index 66) (SNMP ifIndex 417)
          Flags: SNMP-Traps Encapsulation: ENET2
          Input packets : 0
          Output packets: 0
          Protocol eth-switch
          Flags: None

show interfaces user@switch> show interfaces xe-0/1/0 brief
xe-0/1/0 brief Physical interface: xe-0/1/0, Enabled, Physical link is Up
               Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, Loopback: Disabled,
               Source filtering: Disabled, Flow control: Enabled
               Device flags   : Present Running
               Interface flags: SNMP-Traps Internal: 0x0
               Link flags     : None

               Logical interface xe-0/1/0.0
               Flags: SNMP-Traps Encapsulation: ENET2
               eth-switch

show interfaces user@switch> show interfaces xe-4/1/0 detail
xe-4/1/0 detail Physical interface: xe-4/1/0, Enabled, Physical link is Up
               Interface index: 387, SNMP ifIndex: 369, Generation: 390
               Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex,

```

```

BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
Source filtering: Disabled, Flow control: Enabled
Device flags   : Present Running
Interface flags: 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: 00:23:9c:03:8e:70, Hardware address: 00:23:9c:03:8e:70
Last flapped   : 2009-05-12 08:01:04 UTC (00:13:49 ago)
Statistics last cleared: Never
Traffic statistics:
Input bytes   :          4945644          48576 bps
Output bytes  :              0          0 bps
Input packets:          3258          4 pps
Output packets:              0          0 pps
IPv6 transit statistics:
Input bytes   :              0
Output bytes  :              0
Input packets:              0
Output packets:              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 : None
Active defects: None

Logical interface xe-4/1/0.0 (Index 66) (SNMP ifIndex 417) (Generation 158)
Flags: SNMP-Traps Encapsulation: ENET2
Traffic 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
Protocol eth-switch, Generation: 174, Route table: 0
Flags: None
Input Filters: f1,
Output Filters: f2,,,,

```

show interfaces
xe-6/0/39 extensive

```

user@switch> show interfaces xe-6/0/39 extensive
Physical interface: xe-6/0/39, Enabled, Physical link is Up
Interface index: 291, SNMP ifIndex: 1641, Generation: 316
Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex,
BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,

```



```

Source filtering: Disabled, Flow control: Enabled
Device flags   : Present Running
Interface flags: 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:e2:72:f2:88, Hardware address: 00:19:e2:72:f2:88
Last flapped   : 2010-05-13 14:49:43 PDT (1d 00:14 ago)
Statistics last cleared: Never
Traffic statistics:
Input bytes   : 49625962140160      4391057408 bps
Output bytes  : 47686985710805      4258984960 bps
Input packets : 387702829264        4288139 pps
Output packets: 372554570944        4159166 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: 1, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,

FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Ingress queues: 2 supported, 2 in use
Queue counters:   Queued packets  Transmitted packets  Dropped packets
  Low priority      0             336342805223        7986622358
  High priority     0             0
Egress queues: 8 supported, 8 in use
Queue counters:   Queued packets  Transmitted packets  Dropped packets
0 best-effort      0             333760130103        0
1 assured-forw     0             0                  0
2 mcast-be         0             274948977           0
3 queue3           0             0                  0
4 mcast-ef         0             0                  0
5 expedited-fo     0             0                  0
6 mcast-af         0             0                  0
7 network-cont     0             46613              0
PFE Egress queues: 8 supported, 8 in use
Queue counters:   Queued packets  Transmitted packets  Dropped packets
0 best-effort      0             737867061290        5595302082
1 assured-forw     0             0                  0
2 mcast-be         0             0                  0
3 queue3           0             0                  0
4 mcast-ef         0             0                  0
5 expedited-fo     0             0                  0
6 mcast-af         0             0                  0
7 network-cont     0             97800              0
Active alarms : None
Active defects : None
MAC statistics:
Total octets      49625962140160  47686985710805
Total packets     387702829264  372554570944
Unicast packets   387702829264  372554518472
Broadcast packets 0          2
Multicast packets 0          52470
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
Code violations              0
Packet Forwarding Engine configuration:
  Destination slot: 6
CoS information:
  Direction : Output
  CoS transmit queue        Bandwidth      Buffer Priority  Limit
                             %      bps      %      usec
0 best-effort               75    7500000000    75      0      low  none
2 mcast-be                   20    2000000000    20      0      low  none
7 network-cont               5     500000000     5      0      low  none
```

Logical interface xe-6/0/39.0 (Index 1810) (SNMP ifIndex 2238) (Generation 1923)

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

show ipv6 neighbors

Syntax	show ipv6 neighbors
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.3 for EX Series switches.
Description	Display information about the IPv6 neighbor cache.
Options	This command has no options.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> clear ipv6 neighbors on page 202
List of Sample Output	show ipv6 neighbors on page 277 show ipv6 neighbors on page 277
Output Fields	Table 40 on page 277 describes the output fields for the show ipv6 neighbors command. Output fields are listed in the approximate order in which they appear.

Table 40: show ipv6 neighbors Output Fields

Field Name	Field Description
IPv6 Address	Name of the IPv6 interface.
Linklayer Address	Link-layer address.
State	State of the link: up , down , incomplete , reachable , stale , or unreachable .
Exp	Number of seconds until the entry expires.
Rtr	Whether the neighbor is a routing device: yes or no .
Secure	Whether this entry was created using the Secure Neighbor Discovery (SEND) protocol: yes or no .
Interface	Name of the interface.

Sample Output

```

show ipv6 neighbors user@host> show ipv6 neighbors
IPv6 Address      Linklayer Address  State      Exp  Rtr  Interface
fe80::2a0:c9ff:fe5b:4c1e  00:a0:c9:5b:4c:1e  reachable  15   yes  fxp0.0

show ipv6 neighbors user@host > show ipv6 neighbors

```

Interface	IPv6 Address	Linklayer Address	State	Exp	Rtr	Secure
ge-3/2/0.0	fe80::14fb:5dcf:54bd:ff76	00:90:69:a0:a8:bc	stale	1113	yes	yes

show lacp interfaces

Syntax	<code>show lacp interfaces <i>interface-name</i></code>
Release Information	Command introduced in Junos 10.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display Link Aggregation Control Protocol (LACP) information about the specified aggregated Ethernet or Gigabit Ethernet interface.
Options	<p><code>none</code>—Display LACP information for all interfaces.</p> <p><i>interface-name</i>—(Optional) Display LACP information for the specified interface:</p> <ul style="list-style-type: none"> • Aggregated Ethernet—<code>aex</code> • Gigabit Ethernet—<code>ge-fpc/pic/port</code> • 10-Gigabit Ethernet—<code>xe-fpc/pic/port</code>
Required Privilege Level	view
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 21 • 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 27 • Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch • Configuring Aggregated Ethernet Interfaces (CLI Procedure) on page 97 • Configuring Link Aggregation • Configuring Aggregated Ethernet LACP (CLI Procedure) on page 101 • Configuring Aggregated Ethernet LACP • Understanding Aggregated Ethernet Interfaces and LACP on page 8 • Understanding Aggregated Ethernet Interfaces and LACP • Junos OS Interfaces Fundamentals Configuration Guide
List of Sample Output	show lacp interfaces (Aggregated Ethernet) on page 282 show lacp interfaces (QFX Series) on page 282 show lacp interfaces (QFabric Switches) on page 282
Output Fields	Table 41 on page 280 lists the output fields for the <code>show lacp interfaces</code> command. Output fields are listed in the approximate order in which they appear.

Table 41: show lacp interfaces Output Fields

Field Name	Field Description
Aggregated interface	Aggregated Ethernet interface value.
LACP State	<p>LACP state information for each aggregated Ethernet interface:</p> <ul style="list-style-type: none"> For a child interface configured with force-up, LACP state displays FUP along with the interface name. Role—Role played by the interface. It can be one of the following: <ul style="list-style-type: none"> Actor—Local device participating in LACP negotiation. Partner—Remote device participating in LACP negotiation. Exp—Expired state. Yes indicates the actor or partner is in an expired state. No indicates the actor or partner is not in an expired state. Def—Default. Yes indicates that the actor's receive machine is using the default operational partner information, administratively configured for the partner. No indicates the operational partner information in use has been received in an LACP PDU. Dist—Distribution of outgoing frames. No indicates distribution of outgoing frames on the link is currently disabled and is not expected to be enabled. Otherwise, the value is Yes. Col—Collection of incoming frames. Yes indicates collection of incoming frames on the link is currently enabled and is not expected to be disabled. Otherwise, the value is No. Syn—Synchronization. If the value is Yes, the link is considered synchronized. It has been allocated to the correct link aggregation group, the group has been associated with a compatible aggregator, and the identity of the link aggregation group is consistent with the system ID and operational key information transmitted. If the value is No, the link is not synchronized. It is currently not in the right aggregation. Aggr—Ability of aggregation port to aggregate (Yes) or to operate only as an individual link (No). Timeout—LACP timeout preference. Periodic transmissions of LACP PDUs occur at either a slow or fast transmission rate, depending upon the expressed LACP timeout preference (Long Timeout or Short Timeout). Activity—Actor or partner's port activity. Passive indicates the port's preference for not transmitting LAC PDUs unless its partner's control value is Active. Active indicates the port's preference to participate in the protocol regardless of the partner's control value.

Table 41: show lacp interfaces Output Fields (*continued*)

Field Name	Field Description
LACP Protocol	<p>LACP protocol information for each aggregated interface:</p> <ul style="list-style-type: none"> Link state (active or standby) indicated in parentheses next to the interface when link protection is configured. Receive State—One of the following values: <ul style="list-style-type: none"> Current—The state machine receives an LACP PDU and enters the Current state. Defaulted—If no LACP PDU is received before the timer for the Current state expires a second time, the state machine enters the Defaulted state. Expired—If no LACP PDU is received before the timer for the Current state expires once, the state machine enters the Expired state. Initialize—When the physical connectivity of a link changes or a Begin event occurs, the state machine enters the Initialize state. LACP Disabled—If the port is operating in half duplex, the operation of LACP is disabled on the port, forcing the state to LACP Disabled. This state is similar to the Defaulted state, except that the port is forced to operate as an individual port. Port Disabled—If the port becomes inoperable and a Begin event has not occurred, the state machine enters the Port Disabled state. Transmit State—Transmit state of the state machine. One of the following values: <ul style="list-style-type: none"> Fast Periodic—Periodic transmissions are enabled at a fast transmission rate. No Periodic—Periodic transmissions are disabled. Periodic Timer—Transitory state entered when the periodic timer expires. Slow Periodic—Periodic transmissions are enabled at a slow transmission rate. Mux State—State of the multiplexer state machine for the aggregation port. The state is one of the following values: <ul style="list-style-type: none"> Attached—Multiplexer state machine initiates the process of attaching the port to the selected aggregator. Collecting—Yes indicates that the receive function of this link is enabled with respect to its participation in an aggregation. Received frames are passed to the aggregator for collection. No indicates the receive function of this link is not enabled. Collecting Distributing—Collecting and distributing states are merged together to form a combined state (coupled control). Because independent control is not possible, the coupled control state machine does not wait for the partner to signal that collection has started before enabling both collection and distribution. Detached—Process of detaching the port from the aggregator is in progress. Distributing—Yes indicates that the transmit function of this link is enabled with respect to its participation in an aggregation. Frames may be passed down from the aggregator's distribution function for transmission. No indicates the transmit function of this link is not enabled. Waiting—Multiplexer state machine is in a holding process, awaiting an outcome.
LACP Statistics	<p>LACP statistics are returned when the extensive option is used and provides the following information:</p> <ul style="list-style-type: none"> LACP Rx—LACP received counter that increments for each normal hello. LACP Tx—Number of LACP transmit packet errors logged. Unknown Rx—Number of unrecognized packet errors logged. Illegal Rx—Number of invalid packets received.

Sample Output

```

show lacp interfaces (Aggregated Ethernet) user@host> show lacp interfaces ae0 extensive
Aggregated interface: ae0
LACP state:      Role  Exp  Def  Dist  Col  Syn  Aggr  Timeout  Activity

ge-1/0/1FUP      Actor  No   Yes  No   No   No   Yes    Fast    Active
ge-1/0/1FUP      Partner No   Yes  No   No   No   Yes    Fast    Passive
ge-1/0/2         Actor  No   Yes  No   No   No   Yes    Fast    Active
ge-1/0/2         Partner No   Yes  No   No   No   Yes    Fast    Passive

LACP protocol:   Receive State  Transmit State  Mux State
ge-1/0/1FUP      CURRENT        Fast periodic   Collecting
distributing
ge-1/0/2         CURRENT        Fast periodic   Collecting
distributing
ge-1/0/1 (active) CURRENT        Fast periodic   Collecting
distributing
ge-1/0/2 (standby) CURRENT        Fast periodic   WAITING
LACP Statistics: LACP Rx      LACP Tx      Unknown Rx  Illegal Rx
ge-1/0/1         0            0            0           0
ge-1/0/2         0            0            0           0

```

```

show lacp interfaces (QFX Series) user@switch> show lacp interfaces ae0 extensive
Aggregated interface: ae0
LACP state:      Role  Exp  Def  Dist  Col  Syn  Aggr  Timeout  Activity

xe-0/0/1FUP      Actor  No   Yes  No   No   No   Yes    Fast    Active
xe-0/0/1FUP      Partner No   Yes  No   No   No   Yes    Fast    Passive
xe-0/0/2         Actor  No   Yes  No   No   No   Yes    Fast    Active
xe-0/0/2         Partner No   Yes  No   No   No   Yes    Fast    Passive

LACP protocol:   Receive State  Transmit State  Mux State
xe-0/0/1FUP      CURRENT        Fast periodic   Collecting
distributing
xe-0/0/2         CURRENT        Fast periodic   Collecting
distributing
xe-0/0/1 (active) CURRENT        Fast periodic   Collecting
distributing
xe-0/0/2 (standby) CURRENT        Fast periodic   WAITING
LACP Statistics: LACP Rx      LACP Tx      Unknown Rx  Illegal Rx
xe-0/0/1         0            0            0           0
xe-0/0/2         0            0            0           0

```

```

show lacp interfaces (QFabric Switches) 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
LACP Statistics:    LACP Rx    LACP Tx    Unknown Rx    Illegal Rx
node1:xe-0/0/1      0          0          0          0
node2:xe-0/0/2      0          0          0          0

```

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: fe-fpc/pic/port or ge-fpc/pic/port .
Required Privilege Level	view
List of Sample Output	test interface restart-auto-negotiation on page 284
Output Fields	Use the show interfaces extensive 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
--	---