



JUNOS® Software for EX Series Ethernet Switches, Release 10.0: Interfaces

Juniper Networks, Inc.

1194 North Mathilda Avenue
Sunnyvale, California 94089

USA

408-745-2000

www.juniper.net

Revision 1

Published: 2009-11-04

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, the Juniper Networks logo, JUNOS, NetScreen, ScreenOS, and Steel-Belted Radius are registered trademarks of Juniper Networks, Inc. in the United States and other countries. JUNOSe is a trademark 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® Software for EX Series Ethernet Switches, Release 10.0: Interfaces

Copyright © 2009, Juniper Networks, Inc.

All rights reserved. Printed in USA.

Writing:

Editing:

Illustration:

Cover Design:

Revision History

4 November 2009—Revision 1

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

END USER LICENSE AGREEMENT

READ THIS END USER LICENSE AGREEMENT (“AGREEMENT”) BEFORE DOWNLOADING, INSTALLING, OR USING THE SOFTWARE. BY DOWNLOADING, INSTALLING, OR USING THE SOFTWARE OR OTHERWISE EXPRESSING YOUR AGREEMENT TO THE TERMS CONTAINED HEREIN, YOU (AS CUSTOMER OR IF YOU ARE NOT THE CUSTOMER, AS A REPRESENTATIVE/AGENT AUTHORIZED TO BIND THE CUSTOMER) CONSENT TO BE BOUND BY THIS AGREEMENT. IF YOU DO NOT OR CANNOT AGREE TO THE TERMS CONTAINED HEREIN, THEN (A) DO NOT DOWNLOAD, INSTALL, OR USE THE SOFTWARE, AND (B) YOU MAY CONTACT JUNIPER NETWORKS REGARDING LICENSE TERMS.

1. **The Parties.** The parties to this Agreement are (i) Juniper Networks, Inc. (if the Customer’s principal office is located in the Americas) or Juniper Networks (Cayman) Limited (if the Customer’s principal office is located outside the Americas) (such applicable entity being referred to herein as “Juniper”), and (ii) the person or organization that originally purchased from Juniper or an authorized Juniper reseller the applicable license(s) for use of the Software (“Customer”) (collectively, the “Parties”).

2. **The Software.** In this Agreement, “Software” means the program modules and features of the Juniper or Juniper-supplied software, for which Customer has paid the applicable license or support fees to Juniper or an authorized Juniper reseller, or which was embedded by Juniper in equipment which Customer purchased from Juniper or an authorized Juniper reseller. “Software” also includes updates, upgrades and new releases of such software. “Embedded Software” means Software which Juniper has embedded in or loaded onto the Juniper equipment and any updates, upgrades, additions or replacements which are subsequently embedded in or loaded onto the equipment.

3. **License Grant.** Subject to payment of the applicable fees and the limitations and restrictions set forth herein, Juniper grants to Customer a non-exclusive and non-transferable license, without right to sublicense, to use the Software, in executable form only, subject to the following use restrictions:

a. Customer shall use Embedded Software solely as embedded in, and for execution on, Juniper equipment originally purchased by Customer from Juniper or an authorized Juniper reseller.

b. Customer shall use the Software on a single hardware chassis having a single processing unit, or as many chassis or processing units for which Customer has paid the applicable license fees; provided, however, with respect to the Steel-Belted Radius or Odyssey Access Client software only, Customer shall use such Software on a single computer containing a single physical random access memory space and containing any number of processors. Use of the Steel-Belted Radius or IMS AAA software on multiple computers or virtual machines (e.g., Solaris zones) requires multiple licenses, regardless of whether such computers or virtualizations are physically contained on a single chassis.

c. Product purchase documents, paper or electronic user documentation, and/or the particular licenses purchased by Customer may specify limits to Customer’s use of the Software. Such limits may restrict use to a maximum number of seats, registered endpoints, concurrent users, sessions, calls, connections, subscribers, clusters, nodes, realms, devices, links, ports or transactions, or require the purchase of separate licenses to use particular features, functionalities, services, applications, operations, or capabilities, or provide throughput, performance, configuration, bandwidth, interface, processing, temporal, or geographical limits. In addition, such limits may restrict the use of the Software to managing certain kinds of networks or require the Software to be used only in conjunction with other specific Software. Customer’s use of the Software shall be subject to all such limitations and purchase of all applicable licenses.

d. For any trial copy of the Software, Customer’s right to use the Software expires 30 days after download, installation or use of the Software. Customer may operate the Software after the 30-day trial period only if Customer pays for a license to do so. Customer may not extend or create an additional trial period by re-installing the Software after the 30-day trial period.

e. The Global Enterprise Edition of the Steel-Belted Radius software may be used by Customer only to manage access to Customer’s enterprise network. Specifically, service provider customers are expressly prohibited from using the Global Enterprise Edition of the Steel-Belted Radius software to support any commercial network access services.

The foregoing license is not transferable or assignable by Customer. No license is granted herein to any user who did not originally purchase the applicable license(s) for the Software from Juniper or an authorized Juniper reseller.

4. **Use Prohibitions.** Notwithstanding the foregoing, the license provided herein does not permit the Customer to, and Customer agrees not to and shall not: (a) modify, unbundle, reverse engineer, or create derivative works based on the Software; (b) make unauthorized copies of the Software (except as necessary for backup purposes); (c) rent, sell, transfer, or grant any rights in and to any copy of the Software, in any form, to any third party; (d) remove any proprietary notices, labels, or marks on or in any copy of the Software or any product in which the Software is embedded; (e) distribute any copy of the Software to any third party, including as may be embedded in Juniper equipment sold in the secondhand market; (f) use any ‘locked’ or key-restricted feature, function, service, application, operation, or capability without first purchasing the applicable license(s) and obtaining a valid key from Juniper, even if such feature, function, service, application, operation, or capability is enabled without a key; (g) distribute any key for the Software provided by Juniper to any third party; (h) use the Software in any manner that extends or is broader than the uses purchased by Customer from Juniper or an authorized Juniper reseller; (i) use Embedded Software on non-Juniper equipment; (j) use Embedded Software (or make it available for use) on Juniper equipment that the Customer did not originally purchase from Juniper or an authorized Juniper reseller; (k) disclose the results of testing or benchmarking of the Software to any third party without the prior written consent of Juniper; or (l) use the Software in any manner other than as expressly provided herein.

5. **Audit.** Customer shall maintain accurate records as necessary to verify compliance with this Agreement. Upon request by Juniper, Customer shall furnish such records to Juniper and certify its compliance with this Agreement.

6. **Confidentiality.** The Parties agree that aspects of the Software and associated documentation are the confidential property of Juniper. As such, Customer shall exercise all reasonable commercial efforts to maintain the Software and associated documentation in confidence, which at a minimum includes restricting access to the Software to Customer employees and contractors having a need to use the Software for Customer's internal business purposes.
7. **Ownership.** Juniper and Juniper's licensors, respectively, retain ownership of all right, title, and interest (including copyright) in and to the Software, associated documentation, and all copies of the Software. Nothing in this Agreement constitutes a transfer or conveyance of any right, title, or interest in the Software or associated documentation, or a sale of the Software, associated documentation, or copies of the Software.
8. **Warranty, Limitation of Liability, Disclaimer of Warranty.** The warranty applicable to the Software shall be as set forth in the warranty statement that accompanies the Software (the "Warranty Statement"). Nothing in this Agreement shall give rise to any obligation to support the Software. Support services may be purchased separately. Any such support shall be governed by a separate, written support services agreement. TO THE MAXIMUM EXTENT PERMITTED BY LAW, JUNIPER SHALL NOT BE LIABLE FOR ANY LOST PROFITS, LOSS OF DATA, OR COSTS OR PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES, OR FOR ANY SPECIAL, INDIRECT, OR CONSEQUENTIAL DAMAGES ARISING OUT OF THIS AGREEMENT, THE SOFTWARE, OR ANY JUNIPER OR JUNIPER-SUPPLIED SOFTWARE. IN NO EVENT SHALL JUNIPER BE LIABLE FOR DAMAGES ARISING FROM UNAUTHORIZED OR IMPROPER USE OF ANY JUNIPER OR JUNIPER-SUPPLIED SOFTWARE. EXCEPT AS EXPRESSLY PROVIDED IN THE WARRANTY STATEMENT TO THE EXTENT PERMITTED BY LAW, JUNIPER DISCLAIMS ANY AND ALL WARRANTIES IN AND TO THE SOFTWARE (WHETHER EXPRESS, IMPLIED, STATUTORY, OR OTHERWISE), INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NONINFRINGEMENT. IN NO EVENT DOES JUNIPER WARRANT THAT THE SOFTWARE, OR ANY EQUIPMENT OR NETWORK RUNNING THE SOFTWARE, WILL OPERATE WITHOUT ERROR OR INTERRUPTION, OR WILL BE FREE OF VULNERABILITY TO INTRUSION OR ATTACK. In no event shall Juniper's or its suppliers' or licensors' liability to Customer, whether in contract, tort (including negligence), breach of warranty, or otherwise, exceed the price paid by Customer for the Software that gave rise to the claim, or if the Software is embedded in another Juniper product, the price paid by Customer for such other product. Customer acknowledges and agrees that Juniper has set its prices and entered into this Agreement in reliance upon the disclaimers of warranty and the limitations of liability set forth herein, that the same reflect an allocation of risk between the Parties (including the risk that a contract remedy may fail of its essential purpose and cause consequential loss), and that the same form an essential basis of the bargain between the Parties.
9. **Termination.** Any breach of this Agreement or failure by Customer to pay any applicable fees due shall result in automatic termination of the license granted herein. Upon such termination, Customer shall destroy or return to Juniper all copies of the Software and related documentation in Customer's possession or control.
10. **Taxes.** All license fees payable under this agreement are exclusive of tax. Customer shall be responsible for paying Taxes arising from the purchase of the license, or importation or use of the Software. If applicable, valid exemption documentation for each taxing jurisdiction shall be provided to Juniper prior to invoicing, and Customer shall promptly notify Juniper if their exemption is revoked or modified. All payments made by Customer shall be net of any applicable withholding tax. Customer will provide reasonable assistance to Juniper in connection with such withholding taxes by promptly: providing Juniper with valid tax receipts and other required documentation showing Customer's payment of any withholding taxes; completing appropriate applications that would reduce the amount of withholding tax to be paid; and notifying and assisting Juniper in any audit or tax proceeding related to transactions hereunder. Customer shall comply with all applicable tax laws and regulations, and Customer will promptly pay or reimburse Juniper for all costs and damages related to any liability incurred by Juniper as a result of Customer's non-compliance or delay with its responsibilities herein. Customer's obligations under this Section shall survive termination or expiration of this Agreement.
11. **Export.** Customer agrees to comply with all applicable export laws and restrictions and regulations of any United States and any applicable foreign agency or authority, and not to export or re-export the Software or any direct product thereof in violation of any such restrictions, laws or regulations, or without all necessary approvals. Customer shall be liable for any such violations. The version of the Software supplied to Customer may contain encryption or other capabilities restricting Customer's ability to export the Software without an export license.
12. **Commercial Computer Software.** The Software is "commercial computer software" and is provided with restricted rights. Use, duplication, or disclosure by the United States government is subject to restrictions set forth in this Agreement and as provided in DFARS 227.7201 through 227.7202-4, FAR 12.212, FAR 27.405(b)(2), FAR 52.227-19, or FAR 52.227-14(ALT III) as applicable.
13. **Interface Information.** To the extent required by applicable law, and at Customer's written request, Juniper shall provide Customer with the interface information needed to achieve interoperability between the Software and another independently created program, on payment of applicable fee, if any. Customer shall observe strict obligations of confidentiality with respect to such information and shall use such information in compliance with any applicable terms and conditions upon which Juniper makes such information available.
14. **Third Party Software.** Any licensor of Juniper whose software is embedded in the Software and any supplier of Juniper whose products or technology are embedded in (or services are accessed by) the Software shall be a third party beneficiary with respect to this Agreement, and such licensor or vendor shall have the right to enforce this Agreement in its own name as if it were Juniper. In addition, certain third party software may be provided with the Software and is subject to the accompanying license(s), if any, of its respective owner(s). To the extent portions of the Software are distributed under and subject to open source licenses obligating Juniper to make the source code for such portions publicly available (such as the GNU General Public License ("GPL") or the GNU Library General Public License ("LGPL")), Juniper will make such source code portions (including Juniper modifications, as appropriate) available upon request for a period of up to three years from the date of distribution. Such request can be made in writing to Juniper Networks, Inc., 1194 N. Mathilda Ave., Sunnyvale, CA 94089, ATTN: General Counsel. You may obtain a copy of the GPL at <http://www.gnu.org/licenses/gpl.html>, and a copy of the LGPL at <http://www.gnu.org/licenses/lgpl.html>.
15. **Miscellaneous.** This Agreement shall be governed by the laws of the State of California without reference to its conflicts of laws principles. The provisions of the U.N. Convention for the International Sale of Goods shall not apply to this Agreement. For any disputes arising under this Agreement, the Parties hereby consent to the personal and exclusive jurisdiction of, and venue in, the state and federal courts within Santa Clara County, California. This Agreement constitutes the entire and sole agreement between Juniper and the Customer with respect to the Software, and supersedes all prior and contemporaneous

agreements relating to the Software, whether oral or written (including any inconsistent terms contained in a purchase order), except that the terms of a separate written agreement executed by an authorized Juniper representative and Customer shall govern to the extent such terms are inconsistent or conflict with terms contained herein. No modification to this Agreement nor any waiver of any rights hereunder shall be effective unless expressly assented to in writing by the party to be charged. If any portion of this Agreement is held invalid, the Parties agree that such invalidity shall not affect the validity of the remainder of this Agreement. This Agreement and associated documentation has been written in the English language, and the Parties agree that the English version will govern. (For Canada: Les parties aux présentes confirment leur volonté que cette convention de même que tous les documents y compris tout avis qui s'y rattache, soient rédigés en langue anglaise. (Translation: The parties confirm that this Agreement and all related documentation is and will be in the English language)).

Table of Contents

About This Topic Collection	xiii
How to Use This Guide	xiii
List of EX Series Guides for JUNOS Release 10.0	xiii
Downloading Software	xiv
Documentation Symbols Key	xv
Documentation Feedback	xvi
Requesting Technical Support	xvii
Self-Help Online Tools and Resources	xvii
Opening a Case with JTAC	xvii

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	5
Physical Part of an Interface Name	5
Logical Part of an Interface Name	6
Wildcard Characters in Interface Names	7
Understanding Aggregated Ethernet Interfaces and LACP	7
Link Aggregation Group (LAG)	7
Link Aggregation Control Protocol (LACP)	8
Understanding Interface Ranges on EX Series Switches	9
Understanding Layer 3 Subinterfaces	11
Understanding Unicast RPF for EX Series Switches	12
Unicast RPF for EX Series Switches Overview	12
Unicast RPF Implementation for EX Series Switches	13
Global Unicast RPF Implementation	13
Unicast RPF Packet Filtering	13
Bootstrap Protocol (BOOTP) and DHCP Requests	13
Default Route Handling	13
When to Enable Unicast RPF	14
When Not to Enable Unicast RPF	15
ECMP Traffic Handling with Unicast RPF Enabled	16
Understanding IP Directed Broadcast for EX Series Switches	16
IP Directed Broadcast for EX Series Switches Overview	17
IP Directed Broadcast Implementation for EX Series Switches	17

Chapter 4	Verifying Interfaces	79
	Monitoring Interface Status and Traffic	79
	Verifying the Status of a LAG Interface	80
	Verifying That LACP Is Configured Correctly and Bundle Members Are Exchanging LACP Protocol Packets	81
	Verifying the LACP Setup	81
	Verifying That the LACP Packets Are Being Exchanged	81
	Verifying That Layer 3 Subinterfaces Are Working	82
	Verifying Unicast RPF Status	83
	Verifying IP Directed Broadcast Status	85
Chapter 5	Troubleshooting Interfaces	87
	Troubleshooting Network Interfaces on EX3200 and EX4200 Switches	87
	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	87
	The interface on the port in which an SFP or SFP + transceiver is installed in an SFP + uplink module is down	88
	Troubleshooting Uplink Module Installation or Replacement on EX3200 and EX4200 Switches	88
	Virtual Chassis port (VCP) connection does not work	88
	One of the last four network ports on an EX3200 switch with an SFP or SFP + uplink module installed is disabled	89
Chapter 6	Configuration Statements for Interfaces	91
	[edit chassis] Configuration Statement Hierarchy	91
	[edit interfaces] Configuration Statement Hierarchy	91
	802.3ad	93
	aggregated-devices	94
	aggregated-ether-options	95
	auto-negotiation	96
	chassis	97
	description	98
	device-count	99
	duration	100
	ether-options	101
	family ethernet-switching	102
	family ccc	103
	family mpls	103
	filter	104
	force-up	105
	hold-time	105
	inet6-advertise-interval	106
	interface-range	107
	interfaces	108
	lACP	110

Chapter 10	Configuring PoE	193
	Configuring PoE (CLI Procedure)	193
	Configuring PoE (J-Web Procedure)	195
Chapter 11	Verifying PoE Configuration	197
	Monitoring PoE	197
	Verifying Status of PoE Interfaces on an EX Series Switch	198
Chapter 12	Configuration Statements for PoE	199
	[edit poe] Configuration Statement Hierarchy	199
	disable	200
	duration	201
	guard-band	202
	interface	203
	interval	204
	management	205
	maximum-power	206
	priority	207
	telemetries	208
Chapter 13	Operational Mode Commands for PoE	209
	show poe controller	210
	show poe interface	211
	show poe telemetries interface	213

About This Topic Collection

- How to Use This Guide on page xiii
- List of EX Series Guides for JUNOS Release 10.0 on page xiii
- Downloading Software on page xiv
- Documentation Symbols Key on page xv
- Documentation Feedback on page xvi
- Requesting Technical Support on page xvii

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.

The release notes are at

http://www.juniper.net/techpubs/en_US/junos10.0/information-products/topic-collections/release-notes/10.0/junos-release-notes-10.0.pdf.

List of EX Series Guides for JUNOS Release 10.0





Title	Description
<i>Complete Hardware Guide for EX3200 and EX4200 Switches</i>	Component descriptions, site preparation, installation, replacement, and safety and compliance information for EX3200 and EX4200 switches
<i>Complete Hardware Guide for EX8208 Switches</i>	Component descriptions, site preparation, installation, replacement, and safety and compliance information for EX8208 switches
<i>Complete Hardware Guide for EX8216 Switches</i>	Component descriptions, site preparation, installation, replacement, and safety and compliance information for EX8216 switches
<i>Complete Software Guide for JUNOS® Software for EX Series Switches, Release 10.0</i>	Software feature descriptions, configuration examples, and tasks for JUNOS Software for EX Series switches

Title	Description
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</i> .)
<i>JUNOS® Software for EX Series Switches, Release 10.0: Access Control</i>	
<i>JUNOS® Software for EX Series Switches, Release 10.0: Alarms and System Log Messages</i>	
<i>JUNOS® Software for EX Series Switches, Release 10.0: Configuration and File Management</i>	
<i>JUNOS® Software for EX Series Switches, Release 10.0: Class of Service</i>	
<i>JUNOS® Software for EX Series Switches, Release 10.0: Device Security</i>	
<i>JUNOS® Software for EX Series Switches, Release 10.0: Ethernet Switching</i>	
<i>JUNOS® Software for EX Series Switches, Release 10.0: Interfaces</i>	
<i>JUNOS® Software for EX Series Switches, Release 10.0: Layer 3 Protocols</i>	
<i>JUNOS® Software for EX Series Switches, Release 10.0: MPLS</i>	
<i>JUNOS® Software for EX Series Switches, Release 10.0: Multicast</i>	
<i>JUNOS® Software for EX Series Switches, Release 10.0: Network Management and Monitoring</i>	
<i>JUNOS® Software for EX Series Switches, Release 10.0: Port Security</i>	
<i>JUNOS® Software for EX Series Switches, Release 10.0: Routing Policy and Packet Filtering</i>	
<i>JUNOS® Software for EX Series Switches, Release 10.0: Spanning-Tree Protocols</i>	
<i>JUNOS® Software for EX Series Switches, Release 10.0: System Setup</i>	
<i>JUNOS® Software for EX Series Switches, Release 10.0: User and Access Management</i>	
<i>JUNOS® Software for EX Series Switches, Release 10.0: Virtual Systems</i>	

Downloading Software

You can download JUNOS Software 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 <code>configure</code> 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 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 <code>stub</code> 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/customers/support/downloads/710059.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: Interfaces Configuration on page 25
- Configuring Interfaces on page 55
- Verifying Interfaces on page 79
- Troubleshooting Interfaces on page 87
- Configuration Statements for Interfaces on page 91
- Operational Mode Commands for Interfaces on page 133

Chapter 1

Interfaces—Overview

- EX Series Switches Interfaces Overview on page 3
- Understanding Interface Naming Conventions on EX Series Switches on page 5
- Understanding Aggregated Ethernet Interfaces and LACP on page 7
- Understanding Interface Ranges on EX Series Switches on page 9
- Understanding Layer 3 Subinterfaces on page 11
- Understanding Unicast RPF for EX Series Switches on page 12
- Understanding IP Directed Broadcast for EX Series Switches on page 16
- High Availability Features for EX Series Switches Overview on page 18

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 Software Network Interfaces Configuration Guide* at <http://www.juniper.net/techpubs/software/junos/junos100/index.html>.

- Network Interfaces on page 3
- Special Interfaces on page 4

Network Interfaces

Network interfaces connect to the network and carry network traffic. EX Series switches support the following types of network interfaces:

- LAN access interfaces—Use these EX Series ports 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.
- Trunk interfaces—EX Series access switches can be connected to a distribution switch or customer edge (CE) router. 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 to the access switches must also be configured for trunk mode.
- Power over Ethernet (PoE) interfaces—Juniper Networks EX3200 and EX4200 Ethernet Switches provide PoE network ports with the various switch models

providing either 8, 24, or 48 PoE ports. 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.

- 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 *link aggregation group (LAG)* or *bundle*. These aggregated Ethernet interfaces help to balance traffic and increase the uplink bandwidth.

Special Interfaces

Special interfaces include:

- Virtual Chassis port (VCP) interfaces—Each Juniper Networks EX4200 Ethernet Switch has two dedicated *Virtual Chassis ports (VCPs)* on its rear panel. These ports can be used to interconnect two to ten EX4200 Ethernet switches as a *Virtual Chassis*, which functions as a single network entity. See Understanding the High-Speed Interconnection of the Virtual Chassis Members. When you power on EX Series 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. You can also interconnect EX4200 Switches across distances of up to 25 miles (40 km) by using the SFP, SFP + , or XFP uplink module ports. To do so, you must explicitly set the uplink module ports on the members you want to connect as VCPs. See Setting an Uplink Module Port as a Virtual Chassis Port (CLI Procedure). When you set the uplink module ports as uplink VCPs and connect member switches through those uplink VCPs, a LAG is automatically formed when the link speed is the same on connected VCPs and at least two VCPs on one member are connected to at least two VCPs on another member. See Understanding Virtual Chassis Configurations and Link Aggregation.
- Management interface—The Juniper Networks JUNOS Software 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 Software.)
- Virtual management Ethernet (VME) interface— EX4200 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 a Virtual Chassis Configuration.
- 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 EX4200 switches that are configured as a 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 a Virtual Chassis Configuration.

- 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.

Related Topics

- EX3200 and EX4200 Switches Hardware Overview
- EX8208 Switch Hardware Overview
- EX8216 Switch Hardware Overview
- PoE and EX Series Switches Overview on page 181
- Understanding Interface Naming Conventions on EX Series Switches on page 5
- Understanding Aggregated Ethernet Interfaces and LACP on page 7
- Understanding Interface Ranges on EX Series Switches on page 9
- Understanding Layer 3 Subinterfaces on page 11

Understanding Interface Naming Conventions on EX Series Switches

Juniper Networks EX Series Ethernet Switches use a naming convention for defining the interfaces that is similar to that of other platforms running under Juniper Networks JUNOS Software. This topic provides brief information on the naming conventions used for interfaces on EX Series switches. For additional information, see the *JUNOS Software Network Interfaces Configuration Guide* at <http://www.juniper.net/techpubs/software/junos/junos94/index.html>.

- Physical Part of an Interface Name on page 5
- Logical Part of an Interface Name on page 6
- Wildcard Characters in Interface Names on page 7

Physical Part of an Interface Name

Interfaces in JUNOS Software 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
 - **fe**—Fast Ethernet interface
- *fpc*—EX Series interfaces use the following convention for the FPC portion of interface names:

- On Juniper Networks EX3200 Ethernet Switches and standalone Juniper Networks EX4200 Ethernet Switches (not configured in a Virtual Chassis), the FPC number portion is always **0**.
- On EX4200 switches configured in a Virtual Chassis, the FPC number indicates the member number of the switch within the Virtual Chassis, from **0** through **9**.
- On Juniper Networks EX8200 Ethernet Switches, the FPC number indicates the slot number of the line card that contains the physical interface.
- *pic*—EX Series interfaces use the following convention for the PIC (Physical Interface Card) number portion of interface names:
 - On EX3200 and EX4200 switches, the PIC number is **0** for all built-in interfaces (interfaces that are not on an uplink module).
 - On uplink modules in EX3200 and EX4200 switches, the PIC number is **1**.
 - On EX8200 switches, the PIC number is always **0**.
- *port*—EX Series interfaces use the following convention for port numbers:
 - On EX3200 and EX4200 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.
 - On uplink modules in EX3200 and EX4200 switches, ports are labeled from left to right starting with **0**. Uplink modules provide either 2 or 4 ports.
 - On 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.

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*. 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

When you configure aggregated Ethernet interfaces, you configure a logical interface that is called a *bundle* or a *LAG*. Each LAG can include up to eight Ethernet interfaces.

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 Topics**
- EX Series Switches Interfaces Overview on page 3
 - Front Panel of an EX3200 Switch
 - Front Panel of an EX4200 Switch
 - Slot Numbering for an EX8208 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.

LACP, a subcomponent of IEEE 802.3ad, provides additional functionality for LAGs.

- Link Aggregation Group (LAG) on page 7
- Link Aggregation Control Protocol (LACP) on page 8

Link Aggregation Group (LAG)

You configure a LAG by specifying the link number as a physical device and then associating a set of ports with the link. All the ports must have the same speed and be in full-duplex mode. Juniper Networks JUNOS Software for EX Series Ethernet Switches assigns a unique ID and port priority to each port. The ID and priority are not configurable. When configuring LAGs, consider the following guidelines:

- Up to 12 Ethernet interfaces can be grouped to form a LAG.
- Up to 64 LAGs are supported in a Virtual Chassis configuration.
- Up to 256 LAGs are supported on Juniper Networks EX8200 Ethernet Switches.
- 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.
- Link Aggregation Control Protocol (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 Virtual Chassis Configurations and Link Aggregation.



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 Juniper Networks EX4200 Ethernet 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 Topics**
- Understanding Virtual Chassis Configurations and Link Aggregation
 - Understanding Redundant Trunk Links on EX Series Switches

- Example: Configuring Aggregated Ethernet High-Speed Uplinks Between a Virtual Chassis Access Switch and a Virtual Chassis Distribution Switch on page 25
- Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between a Virtual Chassis Access Switch and a Virtual Chassis Distribution Switch on page 32
- *JUNOS Software Network Interfaces Configuration Guide* at <http://www.juniper.net/techpubs/software/junos/junos100/index.html>

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 bpdu-block interface`
- `ethernet-switching-options interfaces`
- `ethernet-switching-options redundant-trunk-group group-name interface`
- `ethernet-switching-options secure-access-port interface`
- `ethernet-switching-options voip interface`
- `poe interface`
- `protocols dot1x authentication interface`
- `protocols gvrp interface`
- `protocols igmp interface`

- protocols igmp-snooping vlan *vlan-name* interface
- protocols isis interface
- protocols link-management peer lmp-control-channel interface
- protocols link-management te-link *name* interface
- protocols lldp interface
- protocols lldp-med interface
- protocols mpls interface
- protocols mstp interface
- protocols mstp msti-*id* interface
- protocols mstp msti-*id* vlan *vlan-id* interface
- protocols oam ethernet link-fault-management interface
- protocols ospf area
- protocols pim interface
- protocols rip group *group-name* neighbor
- protocols ripng group *group-name* neighbor
- protocols router-advertisement interface
- protocols router-discovery interface
- protocols rsvp interface
- protocols sflow interfaces
- protocols stp interface
- protocols vstp vlan *vlan-id* interface
- vlans *vlan-name* interface

Related Topics

- EX Series Switches Interfaces Overview on page 3
- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 65
- Configuring Aggregated Ethernet Interfaces (CLI Procedure) on page 69
- Configuring a Layer 3 Subinterface (CLI Procedure)
- *JUNOS Software Network Interfaces Configuration Guide* at <http://www.juniper.net/techpubs/software/junos/junos96/index.html>
- interface-range

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 Software 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. The JUNOS Software 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 Topics**
- EX Series Switches Interfaces Overview on page 3
 - Example: Configuring Layer 3 Subinterfaces for a Distribution Switch and an Access Switch on page 37
 - *JUNOS Software Network Interfaces Configuration Guide* at <http://www.juniper.net/techpubs/software/junos/junos90/index.html>

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. The switch applies unicast RPF globally to all interfaces. Therefore, you should enable unicast RPF only on switches with all symmetrically routed interfaces. (A symmetrically routed interface is an interface that uses the same route in both directions between the source and the destination.)

This topic covers:

- Unicast RPF for EX Series Switches Overview on page 12
- Unicast RPF Implementation for EX Series Switches on page 13
- When to Enable Unicast RPF on page 14
- When Not to Enable Unicast RPF on page 15
- ECMP Traffic Handling with Unicast RPF Enabled 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.

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 switches with all 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

- Global Unicast RPF Implementation on page 13
- Unicast RPF Packet Filtering on page 13
- Bootstrap Protocol (BOOTP) and DHCP Requests on page 13
- Default Route Handling on page 13

Global Unicast RPF Implementation

The switch implements unicast RPF on a global basis. Unicast RPF is globally disabled by default. You cannot enable unicast RPF on a per-interface basis.

- 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.

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 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 14. Symmetrical routing means that if an interface receives a packet, the switch uses the same interface to send a reply to the packet source (the receiving interface matches the forwarding-table entry for the best return path to the source).

Figure 1: Symmetrically Routed Interfaces



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

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

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



NOTE: Because unicast RPF is enabled globally on the switch, ensure that all interfaces are symmetrically routed before you enable unicast RPF. 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.



TIP:

It is best to enable unicast RPF explicitly on either all interfaces or only one interface:


- 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 globally enabled on the switch. The drawback to this approach is that the switch displays unicast RPF status as enabled only on interfaces on which unicast RPF is explicitly enabled, so even though unicast RPF is enabled on all interfaces, its status does not display as enabled on all interfaces.
 - Enabling unicast RPF explicitly on all interfaces makes it easier to know if unicast RPF is enabled on the switch because every interface shows the correct status. (Only interfaces on which you explicitly enable unicast RPF display unicast RPF as 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 enabled on all interfaces.
-

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 if any switch interfaces are asymmetrically routed because unicast RPF is enabled globally on all interfaces. All switch interfaces must be symmetrically routed for you to enable unicast RPF without the risk of the switch's discarding traffic that you want to forward.

ECMP Traffic Handling with Unicast RPF Enabled

The switch does 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 can result in the switch's discarding packets that you want to forward because the unicast RPF filter does not examine the entire ECMP address block.

- Related Topics**
- Example: Configuring Unicast RPF on an EX Series Switch on page 45
 - Configuring Unicast RPF (CLI Procedure) on page 74
 - Disabling Unicast RPF (CLI Procedure) on page 75

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 17
- When to Enable IP Directed Broadcast on page 17
- When Not to Enable IP Directed Broadcast on page 17

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 Topics**
- Example: Configuring IP Directed Broadcast on an EX Series Switch on page 49
 - Configuring IP Directed Broadcast (CLI Procedure) on page 75

High Availability Features for EX Series Switches Overview

High availability refers to the hardware and software components that provide redundancy and reliability for packet-based communications. This topic covers the following high availability features of Juniper Networks EX Series Ethernet Switches:

- VRRP on page 18
- Graceful Protocol Restart on page 20
- Redundant Routing Engines on page 21
- Graceful Routing Engine Switchover on page 21
- Virtual Chassis Software Upgrade and Failover Features on page 22
- Link Aggregation on page 22
- Additional High Availability Features of EX Series Switches on page 22

VRRP

For Gigabit Ethernet interfaces, 10-Gigabit Ethernet interfaces, and logical interfaces on EX Series switches, you can configure the Virtual Router Redundancy Protocol (VRRP). When VRRP is configured, the switches act as virtual routing platforms. VRRP enables hosts on a LAN to make use of redundant routing platforms on that LAN without requiring more than the static configuration of a single default route on the hosts. The VRRP routing platforms share the IP address corresponding to the default route configured on the hosts. At any time, one of the VRRP routing platforms is the master (active) and the others are backups. If the master routing platform fails, one of the backup routing platforms becomes the new master, providing a virtual default routing platform and enabling traffic on the LAN to be routed without relying on a single routing platform. Using VRRP, a backup EX Series switch can take over a failed default switch within a few seconds. This is done with minimum loss of VRRP traffic and without any interaction with the hosts.

VRRP for IPv6 provides a much faster switchover to an alternate default routing platform than IPv6 Neighbor Discovery (ND) procedures. VRRP for IPv6 does not support authentication-type/key.



NOTE: Do not confuse the VRRP master and backup routing platforms with the master and backup member switches of a Virtual Chassis (VC) configuration. The master and backup members of a Virtual Chassis configuration comprise a single host. In a VRRP topology, one host operates as the master routing platform and another operates as the backup routing platform, as shown in Figure 4 on page 20.

Switches running VRRP dynamically elect master and backup routing platforms. You can also force assignment of master and backup routing platforms using priorities from 1 through 255, with 255 being the highest priority. In VRRP operation, the default master routing platform sends advertisements to backup routing platforms at regular intervals. The default interval is 1 second. If the backup routing platforms do not receive an advertisement for a set period, the backup routing platform with the highest priority takes over as master and begins forwarding packets.



NOTE: Priority 255 cannot be set for routed VLAN interfaces (RVIs).

Figure 3 on page 19 illustrates a basic VRRP topology with EX Series switches. In this example, Switches A, B, and C are running VRRP and together they make up a virtual routing platform. The IP address of this virtual routing platform is 10.10.0.1 (the same address as the physical interface of Switch A).

Figure 3: Basic VRRP on EX Series Switches

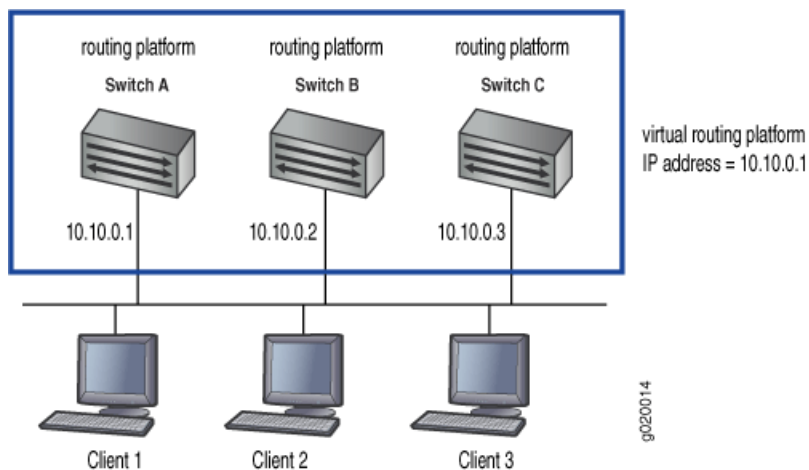
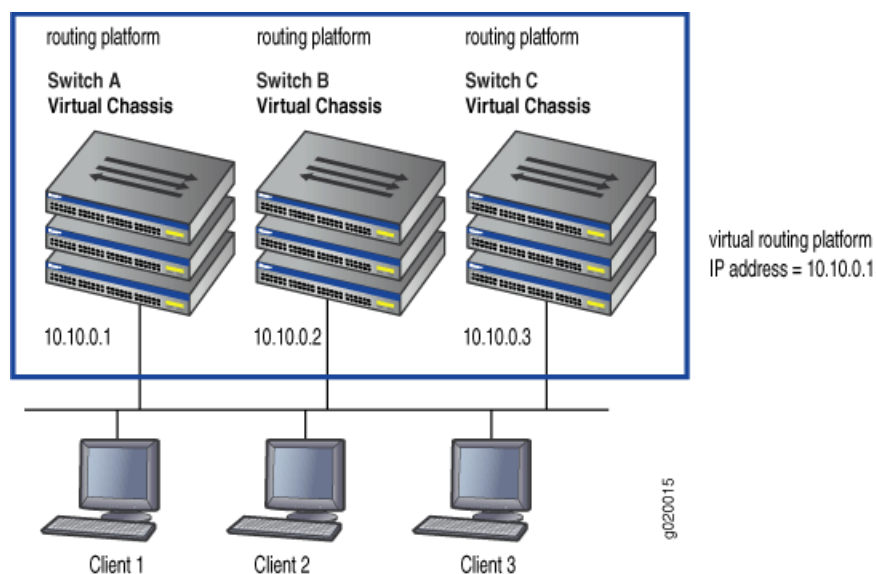


Figure 4 on page 20 illustrates a basic VRRP topology using Virtual Chassis configurations. Switch A, Switch B, and Switch C are each composed of multiple interconnected Juniper Networks EX4200 Ethernet Switches. Each Virtual Chassis configuration operates as a single switch, which is running VRRP, and together they make up a virtual routing platform. The IP address of this virtual routing platform is 10.10.0.1 (the same address as the physical interface of Switch A).

Figure 4: VRRP on Virtual Chassis Switches

Because the virtual routing platform uses the IP address of the physical interface of Switch A, Switch A is the master VRRP routing platform, while switch B and switch C function as backup VRRP routing platforms. Clients 1 through 3 are configured with the default gateway IP address of 10.10.0.1 as the master router, Switch A forwards packets sent to its IP address. If the master routing platform fails, the switch configured with the higher priority becomes the master virtual routing platform and provides uninterrupted service for the LAN hosts. When Switch A recovers, it becomes the master virtual routing platform again.

VRRP is defined in RFC 3768, *Virtual Router Redundancy Protocol*.

Graceful Protocol Restart

With standard implementations of routing protocols, any service interruption requires an affected switch to recalculate adjacencies with neighboring switches, restore routing table entries, and update other protocol-specific information. An unprotected restart of a switch can result in forwarding delays, route flapping, wait times stemming from protocol reconvergence, and even dropped packets. Graceful protocol restart allows a restarting switch and its neighbors to continue forwarding packets without disrupting network performance. Because neighboring switches assist in the restart (these neighbors are called helper switches), the restarting switch can quickly resume full operation without recalculating algorithms from scratch.

On EX Series switches, graceful protocol restart can be applied to aggregate and static routes and for routing protocols (BGP, IS-IS, OSPF, and RIP).

Graceful protocol restart works similarly for the different routing protocols. The main benefits of graceful protocol restart are uninterrupted packet forwarding and temporary suppression of all routing protocol updates. Graceful protocol restart thus allows a switch to pass through intermediate convergence states that are hidden from the rest of the network. Most graceful restart implementations define two types of switches—the restarting switch and the helper switch. The restarting switch requires

rapid restoration of forwarding state information so that it can resume the forwarding of network traffic. The helper switch assists the restarting switch in this process. Individual graceful restart configuration statements typically apply to either the restarting switch or the helper switch.

Redundant Routing Engines

Two to ten EX4200 switches can be interconnected to create a Virtual Chassis configuration that operates as a single network entity. Every Virtual Chassis configuration has a master and a backup. The master acts as the master Routing Engine and the backup acts as the backup Routing Engine. The Routing Engine provides the following functionality:

- Runs various routing protocols
- Provides the forwarding table to the Packet Forwarding Engines (PFEs) in all the member switches of the Virtual Chassis configuration
- Runs other management and control processes for the entire Virtual Chassis configuration

The master Routing Engine, which is in the master of the Virtual Chassis configuration, runs Juniper Networks JUNOS Software in the master role. It receives and transmits routing information, builds and maintains routing tables, communicates with interfaces and Packet Forwarding Engine components of the member switches, and has full control over the Virtual Chassis configuration.

The backup Routing Engine, which is in the backup of the Virtual Chassis configuration, runs JUNOS Software in the backup role. It stays in sync with the master Routing Engine in terms of protocol states, forwarding tables, and so forth. If the master becomes unavailable, the backup Routing Engine takes over the functions that the master Routing Engine performs.

Graceful Routing Engine Switchover

You can configure graceful Routing Engine switchover (GRES) in a Virtual Chassis configuration, allowing the configuration to switch from the master Routing Engine in the master to the backup Routing Engine in the backup with minimal interruption to network communications. When you configure GRES, the backup Routing Engine automatically synchronizes with the master Routing Engine to preserve kernel state information and forwarding state. Any updates to the master Routing Engine are replicated to the backup Routing Engine as soon as they occur. If the kernel on the master Routing Engine stops operating, the master Routing Engine experiences a hardware failure, or the administrator initiates a manual switchover, mastership switches to the backup Routing Engine.

When the backup Routing Engine assumes mastership in a redundant failover configuration (that is, when graceful Routing Engine switchover is not enabled), the Packet Forwarding Engines initialize their state to boot up state before they connect to the new master Routing Engine. In contrast, in a graceful switchover configuration, the Packet Forwarding Engines do not reinitialize their state, but resynchronize their state with the new master Routing Engine. The interruption to the traffic is minimal.

GRES on EX4200 switches supports software features in JUNOS Release 9.2 or later for EX Series switches.

Virtual Chassis Software Upgrade and Failover Features

EX4200 switches provide these features for increased resiliency in Virtual Chassis configurations:

- Virtual Chassis atomic software upgrade—When you upgrade software in a Virtual Chassis configuration, the upgrade will either succeed or fail on all member switches, preventing the situation in which only some of the Virtual Chassis member switches are upgraded.
- Virtual Chassis fast failover—A hardware-assisted failover mechanism that automatically reroutes traffic and reduces traffic loss in the event of a link failure.
- Virtual Chassis split and merge—If there is a disruption to the Virtual Chassis configuration due to member switches failing or being removed from the configuration, the Virtual Chassis configuration splits into two separate Virtual Chassis.

Link Aggregation

You can combine multiple physical Ethernet ports to form a logical point-to-point link, known as a *link aggregation group (LAG)* or *bundle*. A LAG provides more bandwidth than a single Ethernet link can provide. Additionally, link aggregation provides network redundancy by load-balancing traffic across all available links. If one of the links should fail, the system automatically load-balances traffic across all remaining links.

You can select up to eight Ethernet interfaces and include them within a LAG. In an EX4200 Virtual Chassis configuration, the interfaces that form a LAG can be on different members of the Virtual Chassis. See Understanding Virtual Chassis Configurations and Link Aggregation.

Additional High Availability Features of EX Series Switches

To ensure continuous operation, all EX Series switches use field-replaceable power supplies and fan trays. EX3200 and EX4200 switches can use optional field-replaceable uplink modules. EX4200 switches include an option to install a second power supply for redundancy.

The EX4200 switches support connection of Virtual Chassis members using two dedicated Virtual Chassis ports (VCPs) on the rear panel or SFP, SFP + , or XFP uplink module ports. The EX4200 switches also support two internal load-sharing redundant hot-swappable power supplies, field-replaceable fan trays with redundant blowers, and field-replaceable uplink modules that provide SFP, SFP + , or XFP ports.

Notification of hardware issues is provided through system log messages and alarms.

- Related Topics**
- For more information on high availability features, see the *JUNOS Software High Availability Configuration Guide* at <http://www.juniper.net/techpubs/software/junos/junos100/index.html>.
 - Virtual Chassis Overview
 - Understanding Virtual Chassis Components
 - Understanding Virtual Chassis Configurations and Link Aggregation
 - Configuring VRRP for IPv6 (CLI Procedure) on page 76

Chapter 2

Examples: Interfaces Configuration

- Example: Configuring Aggregated Ethernet High-Speed Uplinks Between a Virtual Chassis Access Switch and a Virtual Chassis Distribution Switch on page 25
- Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between a Virtual Chassis Access Switch and a Virtual Chassis Distribution Switch on page 32
- Example: Configuring Layer 3 Subinterfaces for a Distribution Switch and an Access Switch on page 37
- Example: Configuring Unicast RPF on an EX Series Switch on page 45
- Example: Configuring IP Directed Broadcast on an EX Series Switch on page 49

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

EX Series switches allow you to combine one to eight 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.

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

- Requirements on page 25
- Overview and Topology on page 26
- Configuration on page 28
- Verification on page 30
- Troubleshooting on page 31

Requirements

This example uses the following software and hardware components:

- JUNOS 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 Example: Configuring a Virtual Chassis with a Master and Backup in a Single Wiring Closet.
- Configured the uplink ports on the switches as trunk ports. See “Configuring Gigabit Ethernet Interfaces (CLI Procedure)” on page 65.

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 5: Topology for LAGs Connecting a Virtual Chassis Access Switch to a Virtual Chassis Distribution Switch

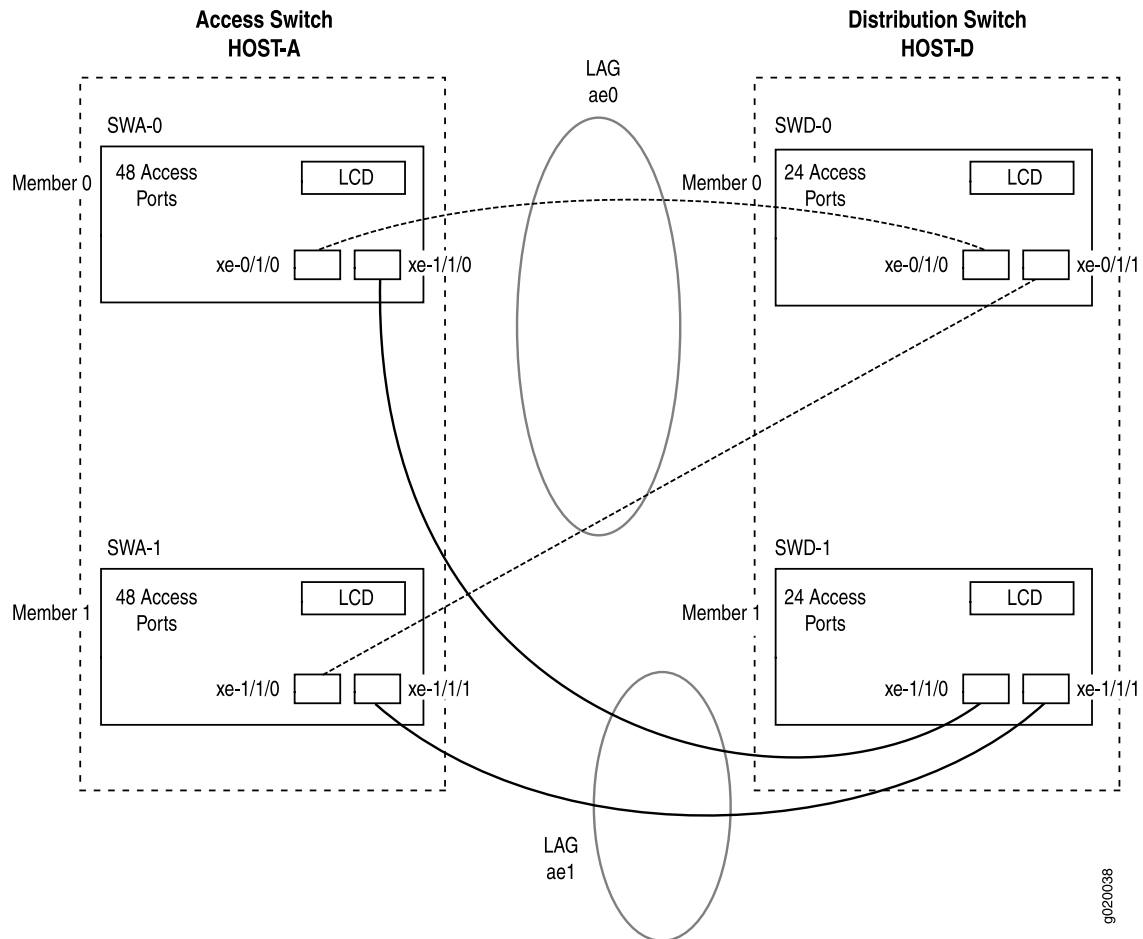


Table 1 details the topology used in this configuration example.

Table 1: Components of the Topology for Connecting Virtual Chassis Access Switches 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

Table 1: Components of the Topology for Connecting Virtual Chassis Access Switches to a Virtual Chassis Distribution Switch (continued)

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

Configuration

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

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

```
[edit]
set chassis aggregated-devices ethernet device-count 2
set interfaces ae0 aggregated-ether-options minimum-links 2
set interfaces ae0 aggregated-ether-options link-speed 10g
set interfaces ae1 aggregated-ether-options minimum-links 2
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 1 family inet address 192.0.2.128/25
set interfaces xe-0/1/0 ether-options 802.ad ae0
set interfaces xe-1/1/0 ether-options 802.ad ae0
set interfaces xe-0/1/1 ether-options 802.ad ae1
set interfaces xe-1/1/1 ether-options 802.ad ae1
```

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

1. Specify the number of LAGs to be created on the chassis:

```
[edit chassis]
user@Host-A# set aggregated-devices ethernet device-count 2
```

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 2
```

3. 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 2
```

- 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
```

- 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.ad ae0
user@Host-A# set xe-1/1/0 ether-options 802.ad ae0
```

- 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.ad ae1
user@Host-A# set xe-1/1/1 ether-options 802.ad ae1
```

- 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
```

- Specify that LAG ae1 belongs to the subnet for the guest broadcast domain:

```
[edit interfaces]
user@Host-A# set ae1 unit 1 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 2;
    }
  }
}
```

```

    unit 0 {
      family inet {
        address 192.0.2.0/25;
      }
    }
  }
  ae1 {
    aggregated-ether-options {
      link-speed 10g;
      minimum-links 2;
    }
    unit 0 {
      family inet {
        address 192.0.2.128/25;
      }
    }
  }
  xe-0/1/0 {
    ether-options {
      802.ad ae0;
    }
  }
  xe-1/1/0 {
    ether-options {
      802.ad ae0;
    }
  }
  xe-0/1/1 {
    ether-options {
      802.ad ae1;
    }
  }
  xe-1/1/1 {
    ether-options {
      802.ad 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 30
- Verifying That LAG ae1 Has Been Created on page 31

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    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.

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

```

Meaning The output shows that the ae1 link is down.

Troubleshooting

Troubleshooting a LAG That 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).

- Related Topics**
- Example: Configuring a Virtual Chassis with a Master and Backup in a Single Wiring Closet
 - Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between a Virtual Chassis Access Switch and a Virtual Chassis Distribution Switch on page 32
 - Example: Connecting an Access Switch to a Distribution Switch.
 - Virtual Chassis Cabling Configuration Examples for EX4200 Switches
 - Installing an Uplink Module in an EX3200 or EX4200 Switch

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

EX Series switches allow you to combine one to eight 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 a Virtual Chassis Access Switch and a Virtual Chassis Distribution Switch” on page 25:

- Requirements on page 32
- Overview and Topology on page 32
- Configuring LACP for the LAGs on the Virtual Chassis Access Switch on page 33
- Configuring LACP for the LAGs on the Virtual Chassis Distribution Switch on page 34
- Verification on page 35
- Troubleshooting on page 36

Requirements

This example uses the following software and hardware components:

- JUNOS 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 Example: Configuring a Virtual Chassis with a Master and Backup in a Single Wiring Closet.
- Configured the uplink ports on the switches as trunk ports. See “Configuring Gigabit Ethernet Interfaces (CLI Procedure)” on page 65.
- Configured the LAGs. See “Example: Configuring Aggregated Ethernet High-Speed Uplinks Between a Virtual Chassis Access Switch and a Virtual Chassis Distribution Switch” on page 25

Overview and Topology

This example assumes that you are already familiar with the Example: Configuring Aggregated Ethernet High-Speed Uplinks between Virtual Chassis Access Switch and Virtual Chassis Distribution Switch. 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. You can configure the interval at which the interfaces send LACP packets by including the periodic statement at the [edit interfaces *interface-name* aggregated-ether-options lACP] hierarchy level.

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 35
- Verifying That the LACP Packets Are Being Exchanged on page 35

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	Co1	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 PDUs are being exchanged.

Troubleshooting

These are some tips for troubleshooting:

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 Topics**
- Example: Connecting an Access Switch to a Distribution Switch
 - Virtual Chassis Cabling Configuration Examples for EX4200 Switches
 - Installing an Uplink Module in an EX3200 or EX4200 Switch

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 37
- Overview and Topology on page 37
- Configuring the Access Switch Subinterfaces on page 38
- Configuring the Distribution Switch Subinterfaces on page 40
- Verification on page 43

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 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 2 on page 38 lists the settings for the example topology.

Table 2: Components of the Topology for Creating Layer 3 Subinterfaces on an Access Switch and a Distribution Switch

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

Configuring the Access Switch Subinterfaces

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

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

Step-by-Step Procedure

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

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

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

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

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

3. Set vlan1's subinterface IP address:

```
[edit interfaces ge-0/1/0]
user@access-switch# set unit 0 family inet address 1.1.1.1/24
```

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

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

5. Set vlan2's subinterface IP address:

```
[edit interfaces ge-0/1/0]
user@access-switch# set unit 1 family inet address 2.1.1.1/24
```

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

```
[edit interfaces ge-0/1/0]
user@access-switch# set unit 2 vlan-id 103
```

7. Set vlan3's subinterface IP address:

```
[edit interfaces ge-0/1/0]
user@access-switch# set unit 2 family inet address 3.1.1.1/24
```

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

```
[edit interfaces ge-0/1/0]
user@access-switch# set unit 3 vlan-id 104
```

9. Set vlan4's subinterface IP address:

```
[edit interfaces ge-0/1/0]
user@access-switch# set unit 3 family inet address 4.1.1.1/24
```

10. Bind vlan5's VLAN ID to the logical interface:

```
[edit interfaces ge-0/1/0]
user@access-switch# set unit 4 vlan-id 105
```

11. Set vlan5's subinterface IP address:

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

Results Check the results of the configuration:

```
user@access-switch> show configuration
interfaces {
  ge-0/1/0 {
    vlan-tagging;
    unit 0 {
      vlan-id 101;
      family inet {
        address 1.1.1.1/24;
      }
    }
    unit 1 {
      vlan-id 102;
      family inet {
        address 2.1.1.1/24;
      }
    }
    unit 2 {
      vlan-id 103;
      family inet {
        address 3.1.1.1/24;
      }
    }
    unit 3 {
      vlan-id 104;
      family inet {
        address 4.1.1.1/24;
      }
    }
    unit 4 {
      vlan-id 105;
      family inet {
        address 5.1.1.1/24;
      }
    }
  }
}
```

Configuring the Distribution Switch Subinterfaces

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

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

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

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

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

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

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

3. Set vlan1's subinterface IP address:

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

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

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

5. Set vlan2's subinterface IP address:

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

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

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

7. Set vlan3's subinterface IP address:

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

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

```
[edit interfaces ge-0/0/0]
user@distribution-switch# set unit 3 vlan-id 104
```

9. Set vlan4's subinterface IP address:

```
[edit interfaces ge-0/0/0]
user@distribution-switch# set unit 3 family inet address 4.1.1.2/24
```

10. Bind vlan5's VLAN ID to the logical interface:

```
[edit interfaces ge-0/0/0]
user@distribution-switch# set unit 4 vlan-id 105
```

11. Set vlan5's subinterface IP address:

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

Results user@distribution-switch> **show configuration**

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

Verification

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

- Verifying That Subinterfaces Were Created on page 43
- Verifying That Traffic Passes Between VLANs on page 43

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-chassis/slot/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 Topics**
- Example: Connecting an Access Switch to a Distribution Switch
 - Configuring a Layer 3 Subinterface (CLI Procedure)

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 global unicast reverse-path forwarding (RPF) on all switch interfaces to filter incoming traffic:

- Requirements on page 45
- Overview and Topology on page 46
- Configuration on page 46
- Verification on page 47

Requirements

This example uses the following software and hardware components:

- JUNOS Release 9.3 or later for EX Series switches
- Two EX3200 switches

Before you configure unicast RPF, make sure that all of the switch interfaces are symmetrically routed (the switch uses the same path in both directions between the source and the destination).

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.

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. All other interfaces on Switch A are also symmetrically routed, because when you enable unicast RPF on any interface, it is thereby enabled globally on all switch interfaces.

The topology of this configuration example uses two EX3200 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 globally on all Switch A interfaces:

CLI Quick Configuration To quickly configure unicast RPF on a switch to help prevent DoS/DDoS attacks, 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 Switch A interfaces to perform unicast RPF filtering:

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 47

Verifying That Unicast RPF Is Enabled on the Switch

Purpose Verify that unicast RPF is enabled.

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

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

```

Input packets:          0
Output packets:        0
Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
  L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
  FIFO errors: 0, Resource errors: 0
Output errors:
  Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,

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

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

Active alarms : LINK
Active defects : LINK
MAC statistics:
  Receive          Transmit
  Total octets    0              0
  Total packets  0              0
  Unicast packets 0              0
  Broadcast packets 0            0
  Multicast packets 0            0
  CRC/Align errors 0              0
  FIFO errors      0              0
  MAC control frames 0              0
  MAC pause frames 0              0
  Oversized frames 0
  Jabber frames    0
  Fragment frames  0
  VLAN tagged frames 0
  Code violations  0
Filter statistics:
  Input packet count 0
  Input packet rejects 0
  Input DA rejects 0
  Input SA rejects 0
  Output packet count 0
  Output packet pad count 0
  Output packet error count 0
  CAM destination filters: 0, CAM source filters: 0
Autonegotiation information:
  Negotiation status: Incomplete
Packet Forwarding Engine configuration:
  Destination slot: 1

Logical interface ge-1/0/10.0 (Index 69) (SNMP ifIndex 59) (Generation 135)
Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
Traffic statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
IPv6 transit statistics:
  Input bytes : 0

```

```

Output bytes : 0
Input packets: 0
Output packets: 0
Local statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Transit statistics:
Input bytes : 0 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. This confirms that unicast RPF is enabled on interface `ge-1/0/10` and thereby on all switch interfaces. Only the interface on which you configured unicast RPF shows the correct unicast RPF configuration status. If you check the unicast RPF status on an interface on which you did not explicitly configure it, the unicast RPF flag is not displayed, even though unicast RPF is implicitly enabled.

- Related Topics**
- Configuring Unicast RPF (CLI Procedure) on page 74
 - Disabling Unicast RPF (CLI Procedure) on page 75

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 49
- Overview and Topology on page 50
- Configuration on page 51

Requirements

This example uses the following software and hardware components:

- JUNOS 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 6 on page 50), 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 6: Topology for IP Directed Broadcast




Table 3 on page 50 shows the settings of the components in this example.

Table 3: 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 13-interface vlan.1
set vlans v0 13-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 13-interface vlan.1
```

```
user@switch# set v0 13-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;
      }
    }
  }
}
```

```
vlan {  
  default;  
  v0 {  
    I3-interface vlan.0;  
  }  
  v1 {  
    I3-interface vlan.1;  
  }  
}
```

Related Topics ■ [Configuring IP Directed Broadcast \(CLI Procedure\) on page 75](#)

Chapter 3

Configuring Interfaces

- Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 55
- Port Role Configuration with the J-Web Interface (with CLI References) on page 61
- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 65
- Setting the Mode on an SFP + Uplink Module (CLI Procedure) on page 68
- Configuring Aggregated Ethernet Interfaces (CLI Procedure) on page 69
- Configuring Aggregated Ethernet Interfaces (J-Web Procedure) on page 70
- Configuring Aggregated Ethernet LACP (CLI Procedure) on page 73
- Configuring Unicast RPF (CLI Procedure) on page 74
- Disabling Unicast RPF (CLI Procedure) on page 75
- Configuring IP Directed Broadcast (CLI Procedure) on page 75
- Configuring VRRP for IPv6 (CLI Procedure) on page 76

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. From the Configure menu, select **Interfaces > Ports**.

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

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.

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.
 - 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 4 on page 57 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 4: 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 5 on page 60 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. <p>You can also select an existing VoIP VLAN configuration or a new VoIP VLAN configuration to be associated with the interface.</p> <p>NOTE: VoIP is not supported on EX8200 switches.</p> <ol style="list-style-type: none"> 2. Click Details to view CLI commands for this role. 3. Click OK.

Table 4: Port Edit Options (continued)

Field	Function	Your Action
Wireless Access Point	<p>Applies the wireless access point role.</p> <p>The interface family is set to 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.
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 5 on page 60 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.
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	<p>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 61 for details on the CLI commands that are associated with each port role.

NOTE: For an EX8200 switch, dynamic ARP inspection and DHCP snooping parameters are not configured.

VLAN Options

Table 4: 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. <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> <p>Click OK.</p>
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, or 1000 Mbps.
Duplex	Specifies the link mode.	Select one: automatic , half-duplex , or full-duplex .
Description	Describes the link. NOTE: If the interface is part of a link aggregation group (LAG), only the option Description is enabled.	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 4: 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 <code>inet</code> family.	<ol style="list-style-type: none"> To specify an IPv4 address, select the check box <code>IPv4 address</code>. Type an IP address — for example: <code>10.10.10.10</code>. Enter the subnet mask or address prefix. For example, 24 bits represents <code>255.255.255.0</code>. Click OK.
IPv6 Address	Specifies an IPv6 address for the interface. NOTE: If the IP address is cleared, the interface still belongs to the <code>inet</code> family.	<ol style="list-style-type: none"> To specify an IPv6 address, select the check box <code>IPv6 address</code>. Type an IP address—for example: <code>2001:ab8:85a3::8a2e:370:7334</code>. Enter the subnet mask or address prefix. Click OK.

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

- Related Topics**
- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 65
 - Monitoring Interface Status and Traffic on page 79
 - EX Series Switches Interfaces Overview on page 3
 - JUNOS CoS for EX Series Switches Overview
 - Understanding Interface Naming Conventions on EX Series Switches on page 5

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 6 on page 61 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 6: 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>> <i>vlan-id</i> <<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</code> <code>vlan members <i>vlan-members</i></code>

Table 6: 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 <code>edge</code> 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 <code>ethernet-switching</code> .	<code>set interfaces <i>interface</i> unit 0 family ethernet-switching port-mode access</code>
Set Port Mode to <code>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>interfaces</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 7 on page 64 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 <code>ethernet-switching</code>	<code>set interfaces <i>interface</i> unit 0 family ethernet-switching port-mode access</code>
Set port mode to <code>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 <code>edge</code> 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>

Table 6: Port Role Configuration Summary (continued)

Configuration Description	CLI Commands
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.	<code>set interfaces <i>interface</i> unit 0 family inet address <i>ipaddress</i></code>
Set IP address on the port.	
Set class-of-service parameters	<code>set class-of-service interfaces <i>interfaces</i> scheduler-map juniper-port-profile-map</code>
SCHEDULER_MAP = juniper-port-profile-map	<code>set class-of-service interfaces <i>interface</i> unit 0 classifiers ieee-802.1 juniper_ieee_classifier</code>
IEEE_CLASSIFIER = juniper-ieee-classifier	<code>set class-of-service interfaces <i>interface</i> unit 0 classifiers dscp juniper-dscp-classifier</code>
DSCP_CLASSIFIER = juniper-dscp-classifier	
Set CoS configuration	Refer Table 7 on page 64 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	<code>set interfaces <i>interface</i> unit 0 family ethernet-switching port-mode trunk</code>
Set port mode to trunk.	
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.	<code>set class-of-service interfaces <i>interfaces</i> scheduler-map juniper-port-profile-map</code>
SCHEDULER_MAP = juniper-port-profile-map	<code>set class-of-service interfaces <i>interface</i> unit 0 classifiers ieee-802.1 juniper_ieee_classifier</code>
IEEE_CLASSIFIER = juniper_ieee_classifier	<code>set class-of-service interfaces <i>interface</i> unit 0 classifiers dscp juniper-dscp-classifier</code>
DSCP_CLASSIFIER = juniper_dscp_classifier	
Set CoS configuration	Refer to Table 7 on page 64 for details.

Table 7 on page 64 lists the CLI commands for the recommended CoS settings that are committed when the CoS configuration is set.

Table 7: 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> <ul style="list-style-type: none"> ■ <code>set class-of-service schedulers assured-scheduler transmit-rate percent 25</code> ■ <code>set class-of-service schedulers strict-priority-scheduler buffer-size percent 25</code> ■ <code>set class-of-service schedulers strict-priority-scheduler priority low</code>

Table 7: Recommended CoS Settings for Port Roles (continued)

CoS Parameter	CLI Command
best-effort-scheduler	The CLI commands are: <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	The classifiers are: <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 Topics	<ul style="list-style-type: none"> ■ Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 55 ■ Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 65

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
- Specifies Spanning Tree Protocol (STP) and Link Layer Discovery Protocol (LLDP)

This topic describes:

- Configuring VLAN Options and Port Mode on page 65
- Configuring the Link Settings on page 66
- Configuring the IP Options on page 67
- Configuring the Interfaces on the Uplink Module in EX3200 and EX4200 Switches on page 67

Configuring VLAN Options and Port Mode

The factory default configuration includes a default VLAN and enables interfaces for the access port mode. Access interfaces typically connect to network devices such as PCs, printers, IP telephones, and IP cameras.

If you are connecting a desktop phone or wireless access point or a security camera to a PoE port, you can configure some parameters for the PoE interface. The PoE interfaces are enabled by default. For detailed information on the PoE settings, see “Configuring PoE (CLI Procedure)” on page 193.

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 you need to configure the logical port as a trunk port. See “Port Role Configuration with the J-Web Interface (with CLI References)” on page 61 for more information about port configuration.

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
```

Configuring the Link Settings

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

- All the 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 (for the EX-UM-2XFP uplink module) 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/1/port ether-options
```



NOTE: An uplink module in an EX Series switch is always PIC 1. The 10-Gigabit Ethernet interface is available only with the EX-UM-2XFP uplink module.

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

- **802.3ad**—Specify an aggregated Ethernet bundle. See “Configuring Aggregated Ethernet Interfaces (CLI Procedure)” on page 69.
- **auto-negotiation**—Enable or disable auto-negotiation of flow control, link mode, and speed.
- **flow-control**—Enable or disable flow control.
- **link-mode**—Specify full-duplex, half-duplex, or automatic.
- **speed**—Specify 10m, 100m, 1g, or autonegotiation.

Configuring the IP Options

To specify an IP address for the logical unit:

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

Configuring the Interfaces on the Uplink Module in EX3200 and EX4200 Switches

By default, the interfaces on the ports on the uplink module installed in EX3200 or EX4200 switches are enabled. You can disable the interfaces on the uplink module using a CLI command.

To disable an interface on the uplink module:

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

where *interface-name* is the name of the interface you want to disable.

If an interface on the uplink module is disabled, you can enable the interface using a CLI command.

To enable an interface on the uplink module:

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

where *interface-name* is the name of the interface you want to enable.

- Related Topics**
- Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 55
 - Monitoring Interface Status and Traffic on page 79
 - show interfaces ge-
 - show interfaces xe-
 - Understanding Interface Naming Conventions on EX Series Switches on page 5
 - Uplink Modules in EX3200 and EX4200 Switches

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. Reboot the switch.

If you commit the configuration but then do not reboot the switch, the new configuration does not take effect—that is, the operating mode of the uplink module is not changed. 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 Topics**
- Uplink Modules in EX3200 and EX4200 Switches
 - Optical Interface Support in EX3200 and EX4200 Switches

Configuring Aggregated Ethernet Interfaces (CLI Procedure)

Use the link aggregation feature to aggregate one or more links to form a virtual link or aggregation group. 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 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.ad ae0
user@switch#set xe-1/1/0 ether-options 802.ad 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 73.

- Related Topics**
- Configuring Aggregated Ethernet Interfaces (J-Web Procedure) on page 70
 - Example: Configuring Aggregated Ethernet High-Speed Uplinks Between a Virtual Chassis Access Switch and a Virtual Chassis Distribution Switch on page 25
 - Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between a Virtual Chassis Access Switch and a Virtual Chassis Distribution Switch on page 32
 - Verifying the Status of a LAG Interface on page 80
 - Understanding Aggregated Ethernet Interfaces and LACP on page 7

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.

2. Click one of the following:
 - **Add**—Creates an aggregated Ethernet interface, or LAG. Enter information as specified in Table 8 on page 71.
 - **Edit**—Modifies a selected LAG.
 - **Aggregation**—Modifies settings for the selected LAG. Enter information as specified in Table 8 on page 71
 - **VLAN**—Specifies VLAN options for the selected LAG. Enter information as specified in Table 9 on page 72.
 - **IP Option**—Specifies IP options for the selected LAG. Enter information as specified in Table 10 on page 72.
 - **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 8: 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	Specifies the mode in which LACP packets are exchanged between the interfaces. The modes are: <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.	To add interfaces to the LAG, select the interfaces and click Add . Click OK . To remove an interface from the LAG, select the interface and click Remove . NOTE: Only interfaces that are configured with the same speed can be selected together for a LAG.
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 9: 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 10: 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 Topics**
- Configuring Aggregated Ethernet Interfaces (CLI Procedure) on page 69
 - Example: Configuring Aggregated Ethernet High-Speed Uplinks Between a Virtual Chassis Access Switch and a Virtual Chassis Distribution Switch on page 25
 - Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between a Virtual Chassis Access Switch and a Virtual Chassis Distribution Switch on page 32
 - Verifying the Status of a LAG Interface on page 80
 - Configuring Aggregated Ethernet LACP (CLI Procedure) on page 73
 - Understanding Aggregated Ethernet Interfaces and LACP on page 7

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 69

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
```

- Related Topics**
- Configuring Aggregated Ethernet Interfaces (CLI Procedure) on page 69
 - Configuring Aggregated Ethernet Interfaces (J-Web Procedure) on page 70
 - Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between a Virtual Chassis Access Switch and a Virtual Chassis Distribution Switch on page 32
 - Example: Configuring Aggregated Ethernet High-Speed Uplinks Between a Virtual Chassis Access Switch and a Virtual Chassis Distribution Switch on page 25
 - Verifying the Status of a LAG Interface on page 80
 - Understanding Aggregated Ethernet Interfaces and LACP on page 7

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 EX Series 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:

- Ensure that all switch interfaces are symmetrically routed before you enable unicast RPF on an interface. A symmetrically routed interface is an interface that uses the same route in both directions between the source and the destination. 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 globally on all switch interfaces, you only need to configure it explicitly on one interface. However, you can configure it explicitly on every interface or only on some interfaces. Regardless of how many interfaces on which you explicitly enable unicast RPF, unicast RPF is implicitly enabled globally after you explicitly configure it on one interface.

We recommend that you enable unicast RPF explicitly on either all interfaces or only one interface, but that you 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 unicast RPF status as enabled only on interfaces on which unicast RPF is explicitly enabled, so even though unicast RPF is enabled on all interfaces, its status does not display as enabled on all interfaces.
- Enabling unicast RPF explicitly on all interfaces makes it easier to know if unicast RPF is enabled on the switch because every interface shows the correct status. (Only interfaces on which you explicitly enable unicast RPF display unicast RPF as 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.

To enable unicast RPF to filter incoming traffic on all switch interfaces by enabling it on one interface:

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

- Related Topics**
- Example: Configuring Unicast RPF on an EX Series Switch on page 45
 - Verifying Unicast RPF Status on page 83
 - Disabling Unicast RPF (CLI Procedure) on page 75
 - Troubleshooting Unicast RPF
 - Understanding Unicast RPF for EX Series Switches on page 12

Disabling Unicast RPF (CLI Procedure)

Unicast reverse-path forwarding (RPF) can help protect your LAN from denial-of-service (DoS) and distributed denial-of-service (DDoS) attacks on untrusted interfaces. Unicast RPF filters traffic with source addresses that do not use the incoming interface as the best return path back to the source. If the network configuration changes so that an interface that has unicast RPF enabled becomes a trusted interface or becomes asymmetrically routed (the interface that receives a packet is not the best return path to the packet's source), you should disable unicast RPF.

To disable unicast RPF on an EX Series switch, you must delete it from every interface on which you explicitly configured it. 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 switch interfaces.

To disable unicast RPF on all switch interfaces, explicitly disable unicast RPF on every interface on which it was explicitly enabled:

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

- Related Topics**
- Example: Configuring Unicast RPF on an EX Series Switch on page 45
 - Verifying Unicast RPF Status on page 83
 - Configuring Unicast RPF (CLI Procedure) on page 74
 - Understanding Unicast RPF for EX Series Switches on page 12

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 v1

user@switch# set ge-0/0/1.0 family ethernet-switching vlan members v1
```

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 v1 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 Topics**
- [Example: Configuring IP Directed Broadcast on an EX Series Switch on page 49](#)
 - [Understanding IP Directed Broadcast for EX Series Switches on page 16](#)

Configuring VRRP for IPv6 (CLI Procedure)

By configuring the Virtual Router Redundancy Protocol (VRRP) on EX Series switches, you can enable hosts on a LAN to make use of redundant routing platforms on that LAN without requiring more than the static configuration of a single default route on the hosts. You can configure VRRP for IPv6 on Gigabit Ethernet, 10-Gigabit Ethernet, and logical interfaces.

To configure VRRP for IPv6:

1. Configure VRRP group support on interfaces:

```
[edit interfaces interface-name unit logical-unit-number family inet6
address address]
user@switch# set vrrp-inet6-group group-id VRRP number virtual-inet6-address address
virtual-link-local-address ipv6-address
```

You must explicitly define a virtual link local address for each VRRP for IPv6 group. Otherwise, when you attempt to commit the configuration, the commit request fails. The virtual link local address must be on the same subnet as the physical interface address.

2. If you want to configure the priority order in which this switch, functioning as a backup router becomes the master router if the master router becomes nonoperational, configure a priority for this switch:

```
[edit interfaces interface-name unit logical-unit-number family inet6
address address vrrp-inet6-group group-id]
user@switch# set VRRP number
```

3. Specify the interval in milliseconds in which the master router sends advertisement packets to the members of the VRRP group:

```
[edit interfaces interface-name unit logical-unit-number family inet6
address address vrrp-inet6-group group-id]
user@switch# set inet6-advertise-interval milliseconds
```

4. By default, a higher-priority backup router preempts a lower-priority master router.

To explicitly enable the master router to be preempted:

- Include the `preempt` statement at the following hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number family inet6
address address vrrp-inet6-group group-id]
user@switch# set preempt
```

To prohibit a higher-priority backup router from preempting a lower priority master router:

- Include the `no-preempt` statement at the following hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number family inet6
address address vrrp-inet6-group group-id]
user@switch# set no-preempt
```

- Related Topics**
- `show vrrp`
 - High Availability Features for EX Series Switches Overview on page 18

Chapter 4

Verifying Interfaces

- Monitoring Interface Status and Traffic on page 79
- Verifying the Status of a LAG Interface on page 80
- Verifying That LACP Is Configured Correctly and Bundle Members Are Exchanging LACP Protocol Packets on page 81
- Verifying That Layer 3 Subinterfaces Are Working on page 82
- Verifying Unicast RPF Status on page 83
- Verifying IP Directed Broadcast Status on page 85

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.

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.

In order 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.
- **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-` or `show interfaces ge-`.
- To view status and statistics for a specific interface, enter `show interfaces xe-interface-name` or `show interfaces ge- interface-name` .
- To view status and traffic statistics for all interfaces, enter `show interfaces xe-detail`, `show interfaces ge- 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 Topics**
- [Configuring Gigabit Ethernet Interfaces \(J-Web Procedure\) on page 55](#)
 - [Configuring Gigabit Ethernet Interfaces \(CLI Procedure\) on page 65](#)

Verifying the Status of a LAG Interface

Purpose Verify that a 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	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 Topics**
- [Configuring Aggregated Ethernet Interfaces \(CLI Procedure\) on page 69](#)
 - [Configuring Aggregated Ethernet Interfaces \(J-Web Procedure\) on page 70](#)

- Example: Configuring Aggregated Ethernet High-Speed Uplinks Between a Virtual Chassis Access Switch and a Virtual Chassis Distribution Switch on page 25

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

To 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 81
2. Verifying That the LACP Packets Are Being Exchanged on page 81

Verifying the LACP Setup

Purpose Verify that the 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.

```
show lacp interfaces xe-0/1/0
```

```
show lacp interfaces xe-0/1/0
```

```
Aggregated interface: ae0
```

LACP state:	Role	Exp	Def	Dist	Co1	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 the 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 Topics**
- Configuring Aggregated Ethernet LACP (CLI Procedure) on page 73
 - Verifying the Status of a LAG Interface on page 80
 - Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between a Virtual Chassis Access Switch and a Virtual Chassis Distribution Switch on page 32

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 ge-chassis/slot/port 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 Topics**
- Configuring a Layer 3 Subinterface (CLI Procedure)
 - Example: Configuring Layer 3 Subinterfaces for a Distribution Switch and an Access Switch on page 37

Verifying Unicast RPF Status

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

Action Use either the `show interfaces ge- extensive` command or the `show interfaces ge- detail` command 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:
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 `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 does not display.



NOTE: The unicast RPF status displays as enabled only on interfaces for which you have explicitly configured unicast RPF. When you enable unicast RPF on one interface, it is automatically enabled on all switch interfaces including LAGs and RVIs. However, the **uRPF** flag does not display on interfaces for which you have not explicitly configured unicast RPF even though unicast RPF is implicitly enabled on those interfaces.

- Related Topics**
- `show interfaces xe-`
 - Example: Configuring Unicast RPF on an EX Series Switch on page 45
 - Configuring Unicast RPF (CLI Procedure) on page 74
 - Disabling Unicast RPF (CLI Procedure) on page 75
 - Troubleshooting Unicast RPF

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 Topics**
- Configuring IP Directed Broadcast (CLI Procedure) on page 75
 - Example: Configuring IP Directed Broadcast on an EX Series Switch on page 49

Chapter 5

Troubleshooting Interfaces

- Troubleshooting Network Interfaces on EX3200 and EX4200 Switches on page 87
- Troubleshooting Uplink Module Installation or Replacement on EX3200 and EX4200 Switches on page 88

Troubleshooting Network Interfaces on EX3200 and EX4200 Switches

This topic provides troubleshooting information for specific problems related to interfaces on EX3200 and EX4200 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 87
- The interface on the port in which an SFP or SFP + transceiver is installed in an SFP + uplink module is down on page 88

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 or 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 68.

- Related Topics**
- Troubleshooting Uplink Module Installation or Replacement on EX3200 and EX4200 Switches on page 88
 - Monitoring Interface Status and Traffic on page 79
 - Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 65
 - Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 55
 - Removing a Transceiver from an EX Series Switch
 - Uplink Modules in EX3200 and EX4200 Switches
 - EX Series Switches Interfaces Overview on page 3

Troubleshooting Uplink Module Installation or Replacement on EX3200 and EX4200 Switches

This topic provides troubleshooting information for specific problems related to uplink module ports on EX3200 and EX4200 switches.

1. Virtual Chassis port (VCP) connection does not work on page 88
2. One of the last four network ports on an EX3200 switch with an SFP or SFP + uplink module installed is disabled on page 89

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 Module Port as a Virtual Chassis Port (CLI Procedure).

One of the last four network ports on an EX3200 switch with an SFP or SFP+ uplink module installed is disabled

Problem One of the last four built-in ports (ge-0/0/20 through ge-0/0/23 on 24-port models or ge-0/0/44 through ge-0/0/47 on 48-port models) of an EX3200 switch with an SFP or SFP + uplink module installed in it is disabled.

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.

- Related Topics**
- Monitoring Interface Status and Traffic on page 79
 - Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 65
 - Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 55
 - Installing an Uplink Module in an EX3200 or EX4200 Switch
 - Removing a Transceiver from an EX Series Switch
 - Uplink Modules in EX3200 and EX4200 Switches
 - Understanding Virtual Chassis Hardware Configuration on an EX4200 Switch

Chapter 6

Configuration Statements for Interfaces

- [edit chassis] Configuration Statement Hierarchy on page 91
- [edit interfaces] Configuration Statement Hierarchy on page 91

[edit chassis] Configuration Statement Hierarchy

```
chassis {
  aggregated-devices {
    ethernet {
      device-count number;
    }
  }
}
```

- Related Topics** ■ *JUNOS Software Hierarchy and RFC Reference* at <http://www.juniper.net/techpubs/software/junos/junos90/index.html>

[edit interfaces] Configuration Statement Hierarchy

```
interfaces {
  aex {
    aggregated-ether-options {
      lacp mode {
        periodic interval;
      }
    }
  }
  ge-chassis/pic/port {
    description text;
    ether-options {
      802.3ad aex;
      auto-negotiation;
      flow-control;
      link-mode mode;
      speed (speed | auto-negotiation) ;
    }
    mtu bytes;
    no-gratuitous-arp-request;
    unit logical-unit-number {
      ( family ccc; |
        family ethernet-switching {
```

```

        filter input filter-name;
        filter output filter-name;
        native-vlan-id vlan-id;
        port-mode mode;
        vlan {
            members [ ( all | names | vlan-ids ) ];
        }
    } |
    family mpls; )
    proxy-arp;
    vlan-id vlan-id-number;
}
vlan-tagging;
}
}

```

- Related Topics**
- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 65
 - Configuring Aggregated Ethernet Interfaces (CLI Procedure) on page 69
 - Configuring a Layer 3 Subinterface (CLI Procedure)
 - EX Series Switches Interfaces Overview on page 3
 - *JUNOS Software Network Interfaces Configuration Guide* at <http://www.juniper.net/techpubs/software/junos/junos96/index.html>

802.3ad

Syntax 802.3ad aex {
 lACP {
 force-up;
 }

Hierarchy Level [edit interfaces *interface-name* ether-options]

Release Information Statement introduced in JUNOS Release 9.0 for EX Series switches.

Description Specify the aggregated Ethernet logical interface number.

Options aex—Aggregated Ethernet logical interface number.

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

- Related Topics**
- Example: Configuring Aggregated Ethernet High-Speed Uplinks Between a Virtual Chassis Access Switch and a Virtual Chassis Distribution Switch on page 25
 - Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between a Virtual Chassis Access Switch and a Virtual Chassis Distribution Switch on page 32
 - Configuring Aggregated Ethernet Interfaces (CLI Procedure) on page 69
 - Configuring Aggregated Ethernet LACP (CLI Procedure) on page 73
 - Understanding Aggregated Ethernet Interfaces and LACP on page 7
 - *JUNOS Software Network Interfaces Configuration Guide* at <http://www.juniper.net/techpubs/software/junos/junos100/index.html>

aggregated-devices

Syntax aggregated-devices {
 ethernet {
 device-count *number*;
 }
 }

Hierarchy Level [edit chassis]

Release Information Statement introduced in JUNOS Release 9.0 for EX Series switches.

Description Configure properties for aggregated devices on the switch.

 The statements are explained separately.

Default Disabled.

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

- Related Topics**
- *JUNOS Network System Basics Configuration Guide* at <http://www.juniper.net/techpubs/software/junos/junos90/index.html>
 - Example: Configuring Aggregated Ethernet High-Speed Uplinks Between a Virtual Chassis Access Switch and a Virtual Chassis Distribution Switch on page 25
 - Configuring Aggregated Ethernet Interfaces (CLI Procedure) on page 69
 - Understanding Aggregated Ethernet Interfaces and LACP on page 7

aggregated-ether-options

Syntax	<pre>aggregated-ether-options { lacp mode { periodic interval; } }</pre>
Hierarchy Level	[edit interfaces aex]
Release Information	Statement introduced in JUNOS Release 9.0 for EX Series switches.
Description	<p>Configure properties specific to a specific aggregated Ethernet interface.</p> <p>The statements are explained separately.</p>
Default	Options are not enabled.
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 Topics	<ul style="list-style-type: none"> ■ Example: Configuring Aggregated Ethernet High-Speed Uplinks Between a Virtual Chassis Access Switch and a Virtual Chassis Distribution Switch on page 25 ■ Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between a Virtual Chassis Access Switch and a Virtual Chassis Distribution Switch on page 32 ■ Configuring Aggregated Ethernet Interfaces (CLI Procedure) on page 69 ■ Configuring Aggregated Ethernet LACP (CLI Procedure) on page 73 ■ Understanding Aggregated Ethernet Interfaces and LACP on page 7 ■ <i>JUNOS Network Interfaces Configuration Guide</i> at http://www.juniper.net/techpubs/software/junos/junos95/index.html

auto-negotiation

Syntax	(auto-negotiation no-auto-negotiation);
Hierarchy Level	[edit interfaces <i>interface-name</i> ether-options]
Release Information	Statement introduced in JUNOS Release 9.0 for EX Series switches.
Description	Explicitly enable or disable autonegotiation. <ul style="list-style-type: none">■ auto-negotiation—Enable autonegotiation.■ no-auto-negotiation—Disable autonegotiation. When autonegotiation is disabled, you must explicitly configure link mode and speed options.
Default	Autonegotiation is automatically enabled. No explicit action is taken after the autonegotiation is complete or if the negotiation fails.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<ul style="list-style-type: none">■ Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 65■ Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 55■ <i>JUNOS Software Network Interfaces Configuration Guide</i> at http://www.juniper.net/techpubs/software/junos/junos95/index.html

chassis

Syntax

```
chassis {
  aggregated-devices {
    ethernet {
      device-count number;
    }
  }
}
```

Hierarchy Level [edit]

Release Information Statement introduced in JUNOS Release 9.0 for EX Series switches.

Description Configure chassis-specific properties. Most standard JUNOS configuration statements are available in the JUNOS for EX Series software. This page lists JUNOS statements that you commonly use when configuring EX Series software as well as statements added to support only EX Series switches.

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 Topics ■ *JUNOS Network System Basics Configuration Guide* at <http://www.juniper.net/techpubs/software/junos/junos90>

description

Syntax	description <i>text</i> ;
Hierarchy Level	[edit interfaces <i>ge-chassis/slot/port</i>]
Release Information	Statement introduced in JUNOS Release 9.0 for EX Series switches.
Description	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 or the switch.
Default	No textual description is configured
Options	<i>text</i> —Text to describe the interface. If the text includes spaces, enclose the entire text in straight quotation marks.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<ul style="list-style-type: none">■ Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 65■ Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 55■ <i>JUNOS Software Network Interfaces Configuration Guide</i> at http://www.juniper.net/techpubs/software/junos/junos95/index.html

device-count

Syntax	device-count <i>number</i> ;
Hierarchy Level	[edit chassis aggregated-devices ethernet]
Release Information	Statement introduced in JUNOS Release 9.0 for EX Series switches. Range updated in JUNOS Release 9.5 for EX Series switches.
Description	Configure the number of aggregated logical devices available to the switch.
Default	There is no default. You must configure a value.
Options	<i>number</i> —Maximum number of Ethernet logical interfaces on the switch. Range: 0 through 127 for EX3200 and EX4200 switches; 0 through 254 for EX8200 switches
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<ul style="list-style-type: none"> ■ Example: Configuring Aggregated Ethernet High-Speed Uplinks Between a Virtual Chassis Access Switch and a Virtual Chassis Distribution Switch on page 25 ■ Configuring Aggregated Ethernet Interfaces (CLI Procedure) on page 69 ■ <i>JUNOS Software Network System Basics Configuration Guide</i> at http://www.juniper.net/techpubs/software/junos/junos95/index.html

duration

Syntax	<code>duration hours;</code>
Hierarchy Level	[edit poe interface (all <i>interface-name</i>) telemetries]
Release Information	Statement introduced in JUNOS Release 9.0 for EX Series switches.
Description	Modify the duration for logging telemetries if you are monitoring the per-port power consumption for PoE interfaces.
Options	<i>hours</i> —Hours the logging continues. Range: 1 through 24 hours Default: 1 hour
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Topics	<ul style="list-style-type: none">■ Example: Configuring PoE Interfaces on an EX Series Switch on page 185■ Example: Configuring PoE Interfaces with Different Priorities on an EX Series Switch on page 188■ Configuring PoE (CLI Procedure) on page 193■ Configuring PoE (J-Web Procedure) on page 195■ PoE and EX Series Switches Overview on page 181

ether-options

Syntax

```
ether-options {
  802.3ad aex {
    lacp (802.3ad) {
      force-up;
    }
  }
  auto-negotiation;
  flow-control;
  link-mode mode;
  speed (speed | auto-negotiation);
}
```

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced in JUNOS Release 9.0 for EX Series switches.

Description Configure `ether-options` properties for a Gigabit Ethernet interface on the EX Series switch.

The remaining statements are explained separately.

Default Enabled.

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

Related Topics

- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 65
- Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 55
-
- EX Series Switches Interfaces Overview on page 3
- *JUNOS Software Network Interfaces Configuration Guide* at <http://www.juniper.net/techpubs/software/junos/junos100/index.html>

family ethernet-switching

Syntax family ethernet-switching {
 filter input *filter-name*;
 filter output *filter-name*;
 native-vlan-id *vlan-id*;
 port-mode *mode*;
 vlan {
 members [(all | *names* | *vlan-ids*)];
 }
 }

Hierarchy Level [edit interfaces *ge-chassis/slot/port* unit *logical-unit-number*]

Release Information Statement introduced in JUNOS Release 9.0 for EX Series switches.

Description Configure Ethernet switching protocol family information for the logical interface.

The remaining statements are explained separately.

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

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

- Related Topics**
- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 65
 - Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 55
 - EX Series Switches Interfaces Overview on page 3
 - *JUNOS Software Network Interfaces Configuration Guide* at <http://www.juniper.net/techpubs/software/junos/junos95/index.html>

family ccc

Syntax	family ccc;
Hierarchy Level	[edit interfaces <i>ge-chassis/slot/port</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced in JUNOS Release 9.5 for EX Series switches.
Description	Configure the logical interface as a circuit cross-connect (CCC).
Default	You must configure a logical interface to be able to use the physical device.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<ul style="list-style-type: none"> ■ Example: Configuring MPLS on EX Series Switches ■ Configuring MPLS on Provider Edge Switches (CLI Procedure) ■ <i>JUNOS Software MPLS Applications Configuration Guide</i> at http://www.juniper.net/techpubs/software/junos/junos95/index.html

family mpls

Syntax	family mpls;
Hierarchy Level	[edit interfaces <i>ge-chassis/slot/port</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced in JUNOS Release 9.5 for EX Series switches.
Description	Configure MPLS protocol family information for the logical interface.
Default	You must configure a logical interface to be able to use the physical device.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<ul style="list-style-type: none"> ■ Example: Configuring MPLS on EX Series Switches ■ Configuring MPLS on Provider Edge Switches (CLI Procedure) ■ Configuring MPLS on Provider Switches (CLI Procedure) ■ <i>JUNOS Software MPLS Applications Configuration Guide</i> at http://www.juniper.net/techpubs/software/junos/junos95/index.html

filter

Syntax	filter (input output) <i>filter-name</i> ;
Hierarchy Level	[edit interfaces <i>ge-chassis/slot/port</i> unit <i>logical-unit-number</i> family ethernet-switching]
Release Information	Statement introduced in JUNOS Release 9.0 for EX Series switches.
Description	Apply a firewall filter to traffic entering the port or Layer 3 interface or exiting the Layer 3 interface.
Default	All incoming traffic is accepted unmodified on the port or Layer 3 interface, and all outgoing traffic is sent unmodified from the port or Layer 3 interface.
Options	<p><i>filter-name</i> —Name of a firewall filter defined in the filter statement.</p> <ul style="list-style-type: none"> ■ input—Apply a firewall filter to traffic entering the port or Layer 3 interface. ■ output—Apply a firewall filter to traffic exiting the Layer 3 interface.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Topics	<ul style="list-style-type: none"> ■ Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches ■ Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 65 ■ Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 55 ■ Configuring Firewall Filters (CLI Procedure) ■ Configuring Firewall Filters (J-Web Procedure) ■ Firewall Filters for EX Series Switches Overview ■ <i>JUNOS Software Network Interfaces Configuration Guide</i> at http://www.juniper.net/techpubs/software/junos/junos95/index.html

force-up

Syntax	force-up;
Hierarchy Level	[edit interfaces <i>interface-name</i> ether-options 802.3ad lacp (802.3ad)]
Release Information	Statement introduced in JUNOS 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 Topics	<ul style="list-style-type: none"> ■ Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 65 ■ Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 55 ■ Understanding Aggregated Ethernet Interfaces and LACP on page 7 ■ <i>JUNOS Software Network Interfaces Configuration Guide</i> at http://www.juniper.net/techpubs/software/junos/junos100/index.html

hold-time

Syntax	hold-time <i>seconds</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> vrrp-inet6-group <i>group-id</i>]
Description	Configure the time in seconds after which a backup router with the highest priority preempts the master router.
Options	<i>seconds</i> —Hold-time period.
Required Privilege Level	Statement introduced in JUNOS Release 10.0 for EX Series switches.
Related Topics	<ul style="list-style-type: none"> ■ Configuring VRRP for IPv6 (CLI Procedure) on page 76 ■ High Availability Features for EX Series Switches Overview on page 18

inet6-advertise-interval

Syntax	inet6-advertise-interval <i>milliseconds</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> vrrp-inet6-group <i>group-id</i>]
Release Information	Statement introduced in JUNOS Release 10.0 for EX Series switches.
Description	Configure the interval between Virtual Router Redundancy Protocol (VRRP) IPv6 advertisement packets.
Options	<i>milliseconds</i> —Interval, in milliseconds, between advertisement packets. Range: 100 to 40,000 milliseconds (ms) Default: 1 second
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<ul style="list-style-type: none">■ Configuring VRRP for IPv6 (CLI Procedure) on page 76■ High Availability Features for EX Series Switches Overview on page 18

interface-range

Syntax `interface-range interface-range name {
ether-options {
802.3ad aex ;
auto-negotiation;
flow-control;
link-mode mode;
speed (speed | auto-negotiation) ;
}
hold-time up milliseconds down milliseconds;
member interface-name;
member-range starting-interface name to ending-interface name;
mtu bytes;
}`

Hierarchy Level [edit interfaces]

Release Information Statement introduced in JUNOS Release 10.0 for EX Series switches.

Description Group interfaces that share a common configuration profile.



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

Options `interface-range-name`—Name of the interface range.



NOTE: You can use regular expressions and wildcards to specify the interfaces in the member-range configuration. Do not use wildcards for interface types.

The remaining statements are explained separately.

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

- Related Topics**
- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 65
 - Understanding Interface Ranges on EX Series Switches on page 9
 - EX Series Switches Interfaces Overview on page 3
 - *JUNOS Software Network Interfaces Configuration Guide* at <http://www.juniper.net/techpubs/software/junos/junos100/index.html>

interfaces

```

Syntax interfaces {
    aex {
        aggregated-ether-options {
            lacp mode {
                periodic interval;
            }
        }
    }
    ge-chassis/slot/port {
        description text;
        ether-options {
            802.3ad aex;
            auto-negotiation;
            flow-control;
            link-mode mode;
            speed (speed | auto-negotiation) ;
        }
        mtu bytes;
        no-gratuitous-arp-request;
        unit logical-unit-number {
            (family ccc; |
            family ethernet-switching {
                filter input filter-name;
                filter output filter-name;
                native-vlan-id vlan-id;
                port-mode mode;
                vlan {
                    members [( all | names | vlan-ids)];
                }
            } |
            family mpls;)
            proxy-arp (restricted | unrestricted);
            vlan-id vlan-id-number;
        }
        vlan-tagging;
    }
    interface-range interface-range name {
        member-range starting-interface name to ending-interface name;
        member interface-name;
        mtu bytes;
        hold-time up milliseconds down milliseconds;
        ether-options {
            802.3ad aex;
            auto-negotiation;
            flow-control;
            link-mode mode;
            speed ( speed | auto-negotiation);
        }
    }
}

```

Hierarchy Level	[edit]
Release Information	Statement introduced in JUNOS Release 9.0 for EX Series switches.
Description	Configure interfaces. Most standard JUNOS configuration statements are available in the JUNOS Software for EX Series switches. This topic lists JUNOS statements that you commonly use when configuring EX Series switches as well as statements added to support EX Series switches only. See the <i>JUNOS Software Network Interfaces Configuration Guide</i> at http://www.juniper.net/techpubs/software/junos/junos100/index.html for additional information.
Options	<p><code>aex</code>—Configure an aggregated Ethernet interface.</p> <p><code>ge-chassis/slot/port</code>—Configure a Gigabit Ethernet interface.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p><code>interface</code>—To view this statement in the configuration.</p> <p><code>interface-control</code>—To add this statement to the configuration.</p>
Related Topics	<ul style="list-style-type: none"> ■ Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 65 ■ Configuring Aggregated Ethernet Interfaces (CLI Procedure) on page 69 ■ Configuring a Layer 3 Subinterface (CLI Procedure) ■ EX Series Switches Interfaces Overview on page 3 ■ <i>JUNOS Software Network Interfaces Configuration Guide</i> at http://www.juniper.net/techpubs/software/junos/junos100/index.html

lACP

Syntax `lACP mode {
 periodic interval;
}`

Hierarchy Level [edit interfaces aex aggregated-ether-options]

Release Information Statement introduced in JUNOS Release 9.0 for EX Series switches.

Description Configure the Link Aggregation Control Protocol (LACP).

Default LACP is not enabled.

Options `mode` —LACP mode:

- `active`—Initiate transmission of LACP packets
- `passive`—Respond to LACP packets

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 Topics**
- Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between a Virtual Chassis Access Switch and a Virtual Chassis Distribution Switch on page 32
 - Configuring Aggregated Ethernet LACP (CLI Procedure) on page 73
 - Configuring Aggregated Ethernet Interfaces (J-Web Procedure) on page 70
 - Understanding Aggregated Ethernet Interfaces and LACP on page 7
 - *JUNOS Software Network Interfaces Configuration Guide* at <http://www.juniper.net/techpubs/software/junos/junos95/index.html>

lACP (802.3ad)

Syntax lACP {
 force-up;
}

Hierarchy Level [edit interfaces *interface-name* ether-options 802.3ad]

Release Information Statement introduced in JUNOS 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 Topics**
- Example: Configuring Aggregated Ethernet High-Speed Uplinks Between a Virtual Chassis Access Switch and a Virtual Chassis Distribution Switch on page 25
 - Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between a Virtual Chassis Access Switch and a Virtual Chassis Distribution Switch on page 32
 - Configuring Aggregated Ethernet Interfaces (CLI Procedure) on page 69
 - Configuring Aggregated Ethernet LACP (CLI Procedure) on page 73
 - Understanding Aggregated Ethernet Interfaces and LACP on page 7
 - *JUNOS Software Network Interfaces Configuration Guide* at <http://www.juniper.net/techpubs/software/junos/junos100/index.html>

link-mode

Syntax	link-mode <i>mode</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> ether-options]
Release Information	Statement introduced in JUNOS Release 9.0 for EX Series switches.
Description	Set the device's link-connection characteristic.
Default	The automatic mode is enabled.
Options	<p><i>mode</i> —Link characteristic:</p> <ul style="list-style-type: none"> ■ full-duplex—Connection is full duplex. ■ half-duplex—Connection is half duplex. ■ automatic—Link mode is negotiated. <p>If no-auto-negotiation is specified in ether-options, you can select only full-duplex or half-duplex. If auto-negotiation is specified in ether-options, you can select any mode.</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 Topics	<ul style="list-style-type: none"> ■ Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 65 ■ Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 55 ■ <i>JUNOS Software Network Interfaces Configuration Guide</i> at http://www.juniper.net/techpubs/software/junos/junos95/index.html

member

Syntax	<code>member interface-name;</code>
Hierarchy Level	[edit interfaces interface-range <i>interface-range-name</i>]
Release Information	Statement introduced in JUNOS 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 Topics	<ul style="list-style-type: none">■ Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 65■ Understanding Interface Ranges on EX Series Switches on page 9■ EX Series Switches Interfaces Overview on page 3■ <i>JUNOS Software Network Interfaces Configuration Guide</i> at http://www.juniper.net/techpubs/software/junos/junos100/index.html

members

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

Hierarchy Level `[edit interfaces ge-chassis/slot/port unit logical-unit-number family ethernet-switching vlan]`

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

Description For trunk interfaces, configure the VLANs for which the interface 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 `vlangs` in your configuration mode command line. Note that only one VLAN is displayed for a VLAN range.

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: Each VLAN that is configured must have a specified VLAN ID when you attempt to commit the configuration; otherwise, the configuration commit fails. Also, `all` cannot be the name of a VLAN on the switch.

`names` —Name of one or more VLANs.

`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`.

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

- Related Topics**
- `show ethernet-switching interfaces`
 - `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 65
 - Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 55
 - Creating a Series of Tagged VLANs (CLI Procedure)
 - Understanding Bridging and VLANs on EX Series Switches

- *JUNOS Software Network Interfaces Configuration Guide* at <http://www.juniper.net/techpubs/software/junos/junos96/index.html>

member-range

Syntax `member-range starting-interface-name to ending-interface-name;`

Hierarchy Level [edit interfaces interface-range *interface-range-name*]

Release Information Statement introduced in JUNOS 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:** *Starting interface-name to ending interface-name*—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 Topics**
- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 65
 - Understanding Interface Ranges on EX Series Switches on page 9
 - EX Series Switches Interfaces Overview on page 3
 - *JUNOS Software Network Interfaces Configuration Guide* at <http://www.juniper.net/techpubs/software/junos/junos100/index.html>

mtu

Syntax mtu bytes;**Hierarchy Level** [edit interfaces *interface-name*]**Release Information** Statement introduced in JUNOS Release 9.0 for EX Series switches.**Description** Specify the maximum transmission unit (MTU) size for the media. Changing the media MTU size causes an interface to be deleted and added again. Keep the following points in mind if you are configuring MTU size for jumbo frames on these special types of interfaces:

- **For LAG interfaces**—Configuring the jumbo MTU size on a link aggregation group (LAG) interface (*aex*) automatically configures the jumbo MTU size on the member links.
- **For RVIs**—Jumbo frames of up to 9216 bytes are supported on the routed VLAN interface (RVI), which is named *vlan*. The RVI functions as a logical router. To route jumbo data packets on the RVI, you must configure the jumbo MTU size on the member physical interfaces of the RVI and not on the RVI itself (the *vlan* interface). However, for jumbo control packets—for example, to ping the RVI with a packet size of 6000 bytes or more—you must explicitly configure the jumbo MTU size on the interface named *vlan* (the RVI).

**CAUTION:** Setting or deleting the jumbo MTU size on the RVI (the *vlan* interface) while the switch is transmitting packets might result in dropped packets.

Default 1514 bytes**Options** bytes—MTU size.
Range: 64 through 9216 bytes
Default: 1514 bytes**Required Privilege Level** interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

- Related Topics**
- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 65
 - Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 55
 - Configuring Routed VLAN Interfaces (CLI Procedure)
 - *JUNOS Software Network Interfaces Configuration Guide* at <http://www.juniper.net/techpubs/software/junos/junos100/index.html>

native-vlan-id

Syntax	native-vlan-id <i>vlan-id</i> ;
Hierarchy Level	[edit interfaces <i>ge-fpc/chassis/port</i> unit 0 family ethernet-switching]
Release Information	Statement introduced in JUNOS 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 Topics	<ul style="list-style-type: none"> ■ show vlans ■ show ethernet-switching interfaces ■ Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 65 ■ <i>JUNOS Software Network Interfaces Configuration Guide</i> at http://www.juniper.net/techpubs/software/junos/junos95/index.html ■ Understanding Bridging and VLANs on EX Series Switches

periodic

Syntax	<code>periodic interval;</code>
Hierarchy Level	[edit interfaces aex aggregated-ether-options lacp]
Release Information	Statement introduced in JUNOS Release 9.0 for EX Series switches.
Description	Configure the interval for periodic transmission of LACP packets.
Default	fast
Options	<i>interval</i> —Interval at which to periodically transmit LACP packets: <ul style="list-style-type: none">■ fast—Transmit packets every second. This is the default.■ slow—Transmit packets every 30 seconds.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<ul style="list-style-type: none">■ Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between a Virtual Chassis Access Switch and a Virtual Chassis Distribution Switch on page 32■ Configuring Aggregated Ethernet LACP (CLI Procedure) on page 73■ Understanding Aggregated Ethernet Interfaces and LACP on page 7■ <i>JUNOS Software Network Interfaces Configuration Guide</i> at http://www.juniper.net/techpubs/software/junos/junos95/index.html

port-mode

Syntax	port-mode <i>mode</i> ;
Hierarchy Level	[edit interfaces <i>ge-chassis/slot/port</i> unit <i>logical-unit-number</i> family ethernet-switching]
Release Information	Statement introduced in JUNOS Release 9.0 for EX Series switches.
Description	Configure whether an interface on the switch operates in access or trunk mode.
Default	All switch interfaces are in access mode.
Options	<p>access—Have the interface operate in access mode. In this mode, the interface can be in a single VLAN only. Access interfaces typically connect to network devices such as PCs, printers, IP telephones, and IP cameras.</p> <p>trunk—Have the interface operate in trunk mode. In this mode, the interface can be in multiple VLANs and can multiplex traffic between different VLANs. Trunk interfaces typically connect to other switches and to routers on the LAN.</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 Topics	<ul style="list-style-type: none"> ■ Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 65 ■ Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 55 ■ <i>JUNOS Software Network Interfaces Configuration Guide</i> at http://www.juniper.net/techpubs/software/junos/junos95/index.html ■ Example: Connecting an Access Switch to a Distribution Switch

preempt

Syntax (preempt | no-preempt) {
 hold-time *seconds*;
}

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family inet6 address *address* vrrp-inet6-group *group-id*]

Release Information Statement introduced in JUNOS Release 10.0 for EX Series switches.

Description Configure whether a backup router can preempt a master router:

- **preempt**—Allow the master router to be preempted.
- **no-preempt**—Prohibit the preemption of the master router.

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 Topics**
- [Configuring VRRP for IPv6 \(CLI Procedure\) on page 76](#)
 - [High Availability Features for EX Series Switches Overview on page 18](#)

priority

Syntax `priority number;`

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family inet address *address* [vrrp-group *group-id*],
edit interfaces *interface-name* unit *logical-unit-number* family inet6 address *address* vrrp-inet6-group *group-id*]

Release Information Statement introduced in JUNOS Release 10.0 for EX Series switches.

Description Configure a switch's priority for becoming the master default routing platform. The routing platform with the highest priority within the group becomes the master.

Options *priority*—Routing platform's priority for being elected to be the master router in the VRRP group. A larger value indicates a higher priority for being elected.
Range: 1 through 255
Default: 100 (for backup routers)



NOTE: Priority 255 cannot be assigned to routed VLAN interfaces (RVIs).

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

- Related Topics**
- Configuring VRRP for IPv6 (CLI Procedure) on page 76
 - High Availability Features for EX Series Switches Overview on page 18

rpf-check

Syntax	rpf-check;
Hierarchy Level	[edit interfaces <i>ge-chassis/slot/port</i> unit <i>logical-unit-number</i> family inet]
Release Information	Statement introduced in JUNOS Release 9.3 for EX Series switches.
Description	Enable a reverse-path forwarding check on unicast traffic (except ECMP packets) on all ingress interfaces.
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 Topics	<ul style="list-style-type: none">■ Example: Configuring Unicast RPF on an EX Series Switch on page 45■ Configuring Unicast RPF (CLI Procedure) on page 74■ Disabling Unicast RPF (CLI Procedure) on page 75■ Understanding Unicast RPF for EX Series Switches on page 12

speed

Syntax	speed (<i>speed</i> auto-negotiation) ;
Hierarchy Level	[edit interfaces <i>interface-name</i> ether-options]
Release Information	Statement introduced in JUNOS 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"> ■ 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 ■ 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.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<ul style="list-style-type: none"> ■ Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 65 ■ Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 55 ■ <i>JUNOS Software Network Interfaces Configuration Guide</i> at http://www.juniper.net/techpubs/software/junos/junos95/index.html

targeted-broadcast

Syntax	targeted-broadcast;
Hierarchy Level	[edit interfacesge-chassis/slot/port unit <i>logical-unit-number</i> family inet]
Release Information	Statement introduced in JUNOS 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 Topics	<ul style="list-style-type: none">■ Example: Configuring IP Directed Broadcast on an EX Series Switch on page 49■ Configuring IP Directed Broadcast (CLI Procedure) on page 75■ Understanding IP Directed Broadcast for EX Series Switches on page 16

unit

Syntax `unit logical-unit-number {`
 `(family ccc; |`
 `family ethernet-switching {`
 `filter input filter-name;`
 `filter output filter-name;`
 `native-vlan-id vlan-id;`
 `port-mode mode;`
 `vlan {`
 `members [(all | names | vlan-ids)];`
 `}`
 `} |`
 `family mpls;)`
 `proxy-arp (restricted | unrestricted);`
 `vlan-id vlan-id-number;`
`}`

Hierarchy Level [edit interfaces *ge-chassis/slot/port*]

Release Information Statement introduced in JUNOS 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 *logical-unit-number*—Number of the logical unit.
Range: 0 through 16,384

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 Topics**
- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 65
 - Configuring Aggregated Ethernet Interfaces (CLI Procedure) on page 69
 - EX Series Switches Interfaces Overview on page 3
 - *JUNOS Software Network Interfaces Configuration Guide* at <http://www.juniper.net/techpubs/software/junos/junos100/index.html>

virtual-inet6-address

Syntax virtual-inet6-address [*addresses*];

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family inet6 address *address* vrrp-inet6-group *group-id*]

Release Information Statement introduced in JUNOS Release 10.0 for EX Series switches.

Description Configure the addresses of the virtual routers in a Virtual Router Redundancy Protocol (VRRP) IPv6 group. You can configure up to eight addresses.



NOTE: The address of an aggregated Ethernet interface (a LAG) or a routed VLAN interface (RVI) cannot be assigned as the virtual router address in a VRRP IPv6 group.

Options *addresses*—Addresses of one or more virtual routers. Do not include a prefix length. If the address is the same as the interface's physical address, the interface becomes the master virtual router for the group.

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 Topics

- Configuring VRRP for IPv6 (CLI Procedure) on page 76
- High Availability Features for EX Series Switches Overview on page 18

virtual-link-local-address

Syntax	virtual-link-local-address <i>ipv6-address</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> vrrp-inet6-group <i>group-id</i>]
Release Information	Statement introduced in JUNOS Release 10.0 for EX Series switches.
Description	Configure a virtual link local address for a Virtual Router Redundancy Protocol (VRRP) IPv6 group. You must explicitly define a virtual link local address for each VRRP IPv6 group. The virtual link local address must be in the same subnet as the physical interface address.
Options	<i>ipv6-address</i> —Virtual link local IPv6 address for VRRP for an IPv6 group.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<ul style="list-style-type: none">■ Configuring VRRP for IPv6 (CLI Procedure) on page 76■ High Availability Features for EX Series Switches Overview on page 18

vlan

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

Hierarchy Level [edit interfaces *ge-chassis/slot/port* unit *logical-unit-number* family ethernet-switching]

Release Information Statement introduced in JUNOS Release 9.0 for EX Series switches.

Description For Gigabit Ethernet and aggregated Ethernet interfaces, bind an 802.1Q VLAN tag ID to a logical interface.

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 Topics**
- `show ethernet-switching interfaces`
 - 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 Network Interfaces Configuration Guide* at <http://www.juniper.net/techpubs/software/junos/junos95/index.html>

vlan-id

Syntax `vlan-id vlan-id-number;`

Hierarchy Level `[edit interfaces ge-chassis/slot/port unit logical-unit-number]`

Release Information Statement introduced in JUNOS Release 9.2 for EX Series switches.

Description For Gigabit Ethernet and aggregated Ethernet interfaces only, 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 `vlan-id-number`—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 Topics ■ `vlan-tagging`

- Example: Configuring Layer 3 Subinterfaces for a Distribution Switch and an Access Switch on page 37
- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 65
- Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 55
- Configuring a Layer 3 Subinterface (CLI Procedure)
- *JUNOS Software Network Interfaces Configuration Guide* at <http://www.juniper.net/techpubs/software/junos/junos95/index.html>

vlan-tagging

Syntax	vlan-tagging;
Hierarchy Level	[edit interfaces <i>ge-chassis/pic/port</i>]
Release Information	Statement introduced in JUNOS Release 9.2 for EX Series switches.
Description	Enable VLAN tagging. The platform will receive and forward single-tag frames with 802.1Q VLAN tags.
Default	VLAN tagging is disabled by default.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<ul style="list-style-type: none">■ vlan-id■ Example: Configuring Layer 3 Subinterfaces for a Distribution Switch and an Access Switch on page 37■ Configuring a Layer 3 Subinterface (CLI Procedure)■ <i>JUNOS Software Network Interfaces Configuration Guide</i> at http://www.juniper.net/techpubs/software/junos/junos95/index.html

vrrp-inet6-group

Syntax `vrrp-inet6-group group-id {
 inet6-advertise-interval milliseconds;
 preempt{
 hold-time seconds;
 }
 VRRP number;
 virtual-inet6-address;
 virtual-link-local-address
 }`

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family inet6 address *address*]

Release Information Statement introduced in JUNOS Release 10.0 for EX Series switches.

Description Configure a Virtual Router Redundancy Protocol (VRRP) IPv6 group.

Options *group-id*—VRRP group identifier. If you enable MAC source address filtering on the interface, you must include the virtual MAC address in the list of source MAC addresses that you specify in the **source-address-filter** statement. MAC addresses ranging from 00:00:5e:00:01:00 through 00:00:5e:00:01:ff are reserved for VRRP, as defined in RFC 3768. The VRRP group number must be the decimal equivalent of the last hexadecimal byte of the virtual MAC address.
Range: 0 through 255

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

Related Topics

- Configuring VRRP for IPv6 (CLI Procedure) on page 76
- High Availability Features for EX Series Switches Overview on page 18

Chapter 7

Operational Mode Commands for Interfaces

show ethernet-switching interfaces

Syntax	show ethernet-switching interfaces <brief detail summary> <interface <i>interface-name</i> >
Release Information	Command introduced in JUNOS Release 9.0 for EX Series switches. In JUNOS Release 9.6 for EX Series switches, the following updates were made: <ul style="list-style-type: none"> ■ Blocking field output updated. ■ The default view updated to include information about 802.1Q-tags. ■ The detail view updated to include information VLAN mapping.
Description	Display information about switched Ethernet interfaces.
Options	none—(Optional) Display brief information for Ethernet switching interfaces. brief detail summary—(Optional) Display the specified level of output. interface <i>interface-name</i> —(Optional) Display Ethernet switching information for a specific interface.
Required Privilege Level	view
Related Topics	<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 135</p> <p>show ethernet-switching interfaces ge-0/0/15 brief on page 136</p> <p>show ethernet-switching interfaces ge-0/0/2 detail (Blocked by RTG rtggroup) on page 136</p> <p>show ethernet-switching interfaces ge-0/0/15 detail (Blocked by STP) on page 136</p> <p>show ethernet-switching interfaces ge-0/0/17 detail (Disabled by bpdu-control) on page 136</p> <p>show ethernet-switching interfaces detail (C-VLAN to S-VLAN Mapping) on page 136</p>
Output Fields	Table 11 on page 134 lists the output fields for the show ethernet-switching interfaces command. Output fields are listed in the approximate order in which they appear.

Table 11: show ethernet-switching interfaces Output Fields

Field Name	Field Description	Level of Output
Interface	Name of a switching interface.	All levels
State	Interface state. Values are up and down.	none, brief, detail, summary

Table 11: show ethernet-switching interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
VLAN members	Name of a VLAN.	none, brief, detail, summary
Tag	Number of the 802.1Q-tag.	All levels
Tagging	Specifies whether the interface forwards 802.1Q-tagged or untagged traffic.	All levels
Blocking	The forwarding state of the interface: <ul style="list-style-type: none"> ■ unblocked—Traffic is forwarded on the interface. ■ blocked—Traffic is not being forwarded on the interface. ■ Disabled by bpd control—The interface is disabled due to receiving BPDUs on a protected interface. If the <code>disable-timeout</code> 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 limiting 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 limiting 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, summary
Index	The VLAN index internal to JUNOS Software.	detail
mapping	The C-VLAN to S-VLAN mapping information: <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. 	detail

show ethernet-switching interfacesuser@switch> **show ethernet-switching interfaces**

```

Interface   State  VLAN members      Tag  Tagging  Blocking
-----
ae0.0      up    default          300  untagged unblocked
ge-0/0/2.0 up    vlan300          300  untagged blocked by RTG (rtggroup)
ge-0/0/3.0 up    default          300  untagged blocked by STP
ge-0/0/4.0 down  default          300  untagged MAC limit exceeded
ge-0/0/5.0 down  default          300  untagged MAC move limit exceeded
ge-0/0/6.0 down  default          300  untagged 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 ethernet-switching user@switch> show ethernet-switching interfaces ge-0/0/15 brief
interfaces ge-0/0/15 Interface State VLAN members Tag Tagging Blocking
brief
ge-0/0/15.0 up vlan100 100 tagged blocked by STP
                vlan200 200 tagged blocked by STP

```

```

show ethernet-switching user@switch> show ethernet-switching interfaces ge-0/0/2 detail
interfaces ge-0/0/2 Interface: ge-0/0/2.0, Index: 65, State: up, Port mode: Access
detail (Blocked by RTG VLAN membership:
rtggroup)      vlan300, 802.1Q Tag: 300, untagged, msti-id: 0, blocked by RTG(rtggroup)
Number of MACs learned on IFL: 0

```

```

show ethernet-switching user@switch> show ethernet-switching interfaces ge-0/0/15 detail
interfaces ge-0/0/15 Interface: ge-0/0/15.0, Index: 70, State: up, Port mode: Trunk
detail (Blocked by STP) 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 user@switch> show ethernet-switching interfaces ge-0/0/17 detail
interfaces ge-0/0/17 Interface: ge-0/0/17.0, Index: 71, State: down, Port mode: Trunk
detail (Disabled by VLAN membership:
bpdu-control)      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 user@switch> show ethernet-switching interfaces ge-0/0/6.0 detail
interfaces detail Interface: ge-0/0/6.0, Index: 73, State: up, Port mode: Access
(C-VLAN to S-VLAN VLAN membership:
Mapping)      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 interfaces diagnostics optics

Syntax	show interfaces diagnostics optics <i>interface-name</i>
Release Information	Command introduced in JUNOS Release 10.0 for EX Series switches.
Description	<p>Display diagnostics data and alarms for a Gigabit Ethernet SFP, SFP + , or XFP transceiver installed in an EX3200 and EX4200 switch or a Gigabit Ethernet SFP or SFP + transceiver installed in the line cards installed in an EX8200 switch. 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: <i>ge-fpc/pic/port</i> or <i>xe-fpc/pic/port</i> .
Required Privilege Level	view
Related Topics	<ul style="list-style-type: none"> ■ Monitoring Interface Status and Traffic on page 79 ■ Installing a Transceiver in an EX Series Switch ■ Removing a Transceiver from an EX Series Switch ■ <i>JUNOS Software Network Interfaces Configuration Guide</i> at http://www.juniper.net/techpubs/software/junos/junos100/index.html
List of Sample Output	<p>show interfaces diagnostics optics ge-0/1/0 (SFP Transceiver) on page 140</p> <p>show interfaces diagnostics optics xe-0/1/0 (SFP+ Transceiver) on page 141</p> <p>show interfaces diagnostics optics xe-0/1/0 (XFP Transceiver) on page 142</p>
Output Fields	Table 12 on page 137 lists the output fields for the <code>show interfaces diagnostics optics</code> command. Output fields are listed in the approximate order in which they appear.

Table 12: 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 12: 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.
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.

Table 12: show interfaces diagnostics optics Output Fields (continued)

Field Name	Field Description
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 .
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.

Table 12: show interfaces diagnostics optics Output Fields (continued)

Field Name	Field Description
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.
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.

```

show interfaces user@host> show interfaces diagnostics optics ge-0/1/0
diagnostics optics Physical interface: ge-0/1/0
                        Laser bias current                : 5.444 mA

```

```

ge-0/1/0
(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

```

```

show interfaces
diagnostics optics
xe-0/1/0
(SFP+ Transceiver)
user@host> show interfaces diagnostics optics xe-0/1/0
Physical interface: xe-0/1/0
Laser bias current           : 4.968 mA
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

```

**show interfaces
diagnostics optics
xe-0/1/0
(XFP Transceiver)**

```

user@host> show interfaces diagnostics optics xe-0/1/0
Physical interface: xe-0/1/0
Laser bias current           : 8.029 mA
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 <i>ge-fpc/pic/port</i> <brief detail extensive terse> <descriptions> <media> <snmp-index <i>snmp-index</i> > <statistics>
Release Information	Command introduced in JUNOS Release 9.0 for EX Series switches.
Description	Display status information about the specified Gigabit Ethernet interface.
Options	<p><i>ge-fpc/pic/port</i> —Display standard information about the specified Gigabit 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>snmp-index <i>snmp-index</i> —(Optional) Display information for the specified SNMP index of the interface.</p> <p>statistics—(Optional) Display static interface statistics.</p>
Required Privilege Level	view
Related Topics	<ul style="list-style-type: none"> ■ Monitoring Interface Status and Traffic on page 79 ■ Troubleshooting Network Interfaces on EX3200 and EX4200 Switches on page 87 ■ Troubleshooting an Aggregated Ethernet Interface ■ <i>JUNOS Software Network Interfaces Configuration Guide</i> at http://www.juniper.net/techpubs/software/junos/junos95/index.html
List of Sample Output	<p>show interfaces ge-0/0/0 on page 150</p> <p>show interfaces ge-0/0/0 brief on page 151</p> <p>show interfaces ge-0/0/0 detail on page 151</p> <p>show interfaces ge-0/0/4 extensive on page 152</p>
Output Fields	Table 13 on page 144 lists the output fields for the <code>show interfaces ge-</code> command. Output fields are listed in the approximate order in which they appear.

Table 13: show interfaces ge- Output Fields

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

Table 13: 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 at which the interface is running.	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
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

Table 13: show interfaces ge- 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> <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 this platform.</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 Software 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 13: 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 router 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 craft interface. 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 13: 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 13: 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"> ■ Incomplete—Ethernet interface has the speed or link mode configured. ■ No autonegotiation—Remote Ethernet interface has the speed or link mode configured, or does not perform autonegotiation. ■ Complete—Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful. ■ Link partner status—OK when Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful. ■ 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. ■ 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 both PAUSE on receive and transmit or only PAUSE 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—Information from the link partner: <ul style="list-style-type: none"> ■ 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 both PAUSE on receive and transmit or only PAUSE receive). ■ 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. 	extensive
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

Table 13: show interfaces ge- Output Fields (continued)

Field Name	Field Description	Level of Output
Encapsulation	Encapsulation on the logical interface.	All levels
Protocol	Protocol family.	detail extensive none
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 statics 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	Information about protocol family flags. If unicast Reverse Path Forwarding (uRPF) is explicitly configured on the specified interface, the uRPF flag displays. If uRPF 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 does not display even though uRPF is enabled.	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 address flag.	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 interlace.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive

```

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
ge-0/0/0 brief**

```

user@switch> show interfaces 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 i
ngress 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
ge-0/0/0 detail**

```

user@switch> show interfaces 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 ge-0/0/4 extensive

```

user@switch> show interfaces 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:
  Receive          Transmit
Total octets      0          2989761
Total packets    0          24307
Unicast packets  0           0
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           0
Jabber frames    0           0
Fragment frames  0           0
Code violations  0           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
  %          bps      %          usec
  0 best-effort          95          950000000  95          NA          low
none
  7 network-control     5           50000000  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 xe-

Syntax	show interfaces <i>xe-fpc/pic/port</i> <brief detail extensive terse> <descriptions> <media> <snmp-index <i>snmp-index</i> > <statistics>
Release Information	Command introduced in JUNOS Release 9.0 for EX Series switches.
Description	Display status information about the specified 10-Gigabit Ethernet interface.
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>descriptions—(Optional) Display interface description strings.</p> <p>media—(Optional) Display media-specific information about network interfaces.</p> <p>snmp-index <i>snmp-index</i> —(Optional) Display information for the specified SNMP index of the interface.</p> <p>statistics—(Optional) Display static interface statistics.</p>
Required Privilege Level	view
Related Topics	<ul style="list-style-type: none"> ■ Monitoring Interface Status and Traffic on page 79 ■ Troubleshooting Network Interfaces on EX3200 and EX4200 Switches on page 87 ■ Troubleshooting an Aggregated Ethernet Interface ■ <i>JUNOS Software Network Interfaces Configuration Guide</i> at http://www.juniper.net/techpubs/software/junos/junos95/index.html
List of Sample Output	<p>show interfaces xe-0/1/0 on page 162</p> <p>show interfaces xe-4/1/0 on page 162</p> <p>show interfaces xe-0/1/0 brief on page 163</p> <p>show interfaces xe-4/1/0 detail on page 163</p> <p>show interfaces xe-4/1/0 extensive on page 164</p>
Output Fields	Table 14 on page 156 lists the output fields for the <code>show interfaces xe-</code> command. Output fields are listed in the approximate order in which they appear.

Table 14: show interfaces xe- 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.	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
Link-level type	Encapsulation being used on the physical interface.	All levels
MTU	Maximum transmission unit size on the physical interface.	All levels
Speed	Speed at which the interface is running.	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
LAN-PHY mode	10-Gigabit Ethernet interface operating in Local Area Network Physical Layer Device (LAN PHY) mode. LAN PHY allows 10-Gigabit Ethernet wide area links to use existing Ethernet applications.	All levels
Unidirectional	Unidirectional link mode status for 10-Gigabit Ethernet interface: Enabled or Disabled for parent interface; Rx-only or Tx-only for child interfaces.	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
Wavelength	Configured wavelength, in nanometers (nm).	All levels
Frequency	Frequency associated with the configured wavelength, in terahertz (THz).	All levels
CoS queues	Number of CoS queues configured.	detail extensive none
Schedulers	Number of CoS schedulers configured.	extensive
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	detail extensive

Table 14: show interfaces xe- Output Fields (continued)

Field Name	Field Description	Level of Output
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>Last flapped: year-month-day hour: :minute:second:timezone (hour:minute:second ago)</i> . For example, <i>Last flapped: 2008-01-16 10:52:40 UTC (3d 22:58 ago)</i> .	detail extensive none
Input Rate	Input rate in bits per second (bps) and packets per second (pps).	None specified
Output Rate	Output rate in bps and pps.	None specified
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 this platform.</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 Software 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 by if you configure the <code>ignore-l3-incompletes</code> 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 14: show interfaces xe- 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 router 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
Ingress queues	Total number of ingress queues supported on the specified interface.	extensive
Queue counters (Ingress)	<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. 	extensive

Table 14: 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 router configuration, an alarm can ring the red or yellow alarm bell on the router, or turn on the red or yellow alarm LED on the craft interface. 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
PCS statistics	Physical Coding Sublayer (PCS) fault conditions from the LAN PHY device.	detail extensive
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. ■ VLAN tagged frames—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not. This counter is not supported on EX Series switches and is always displayed as 0. ■ 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 14: show interfaces xe- 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"> ■ Incomplete—Ethernet interface has the speed or link mode configured. ■ No autonegotiation—Remote Ethernet interface has the speed or link mode configured, or does not perform autonegotiation. ■ Complete—Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful. ■ Link partner status—OK when Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful. ■ 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. ■ Flow control—Types of flow control supported by the remote Ethernet device. For Fast Ethernet interfaces, the type is None. 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 both PAUSE on receive and transmit or only PAUSE 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. ■ Local resolution—Information from the link partner: <ul style="list-style-type: none"> ■ 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 both PAUSE on receive and transmit or only PAUSE receive). ■ 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

Table 14: show interfaces xe- Output Fields (continued)

Field Name	Field Description	Level of Output
Packet Forwarding Engine configuration	Information about the configuration of the Packet Forwarding Engine: <ul style="list-style-type: none"> ■ Destination slot—FPC slot number. ■ CoS transmit queue—Queue number and its associated user-configured forwarding class name. ■ Bandwidth %—Percentage of bandwidth allocated to the queue. ■ Bandwidth bps—Bandwidth allocated to the queue (in bps). ■ Buffer %—Percentage of buffer space allocated to the queue. ■ Buffer usec—Amount of 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
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
Protocol	Protocol family.	detail extensive none
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 statics 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

Table 14: show interfaces xe- Output Fields (continued)

Field Name	Field Description	Level of Output
Output Filters	Names of any output filters applied to this interface.	detail extensive
Flags	Information about protocol family flags. If unicast Reverse Path Forwarding (uRPF) is explicitly configured on the specified interface, the uRPF flag displays. If uRPF 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 does not display even though uRPF 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 address flag.	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 interlace.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive

```

show interfaces user@switch> show interfaces xe-0/1/0
xe-0/1/0 Physical interface: xe-0/1/0, Enabled, Physical link is Up
  Interface index: 153, SNMP ifIndex: 69
  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
  CoS queues    : 8 supported, 8 maximum usable queues
  Current address: 00:19:e2:50:c8:99, Hardware address: 00:19:e2:50:c8:99
  Last flapped  : 2008-02-25 05:28:08 UTC (00:12:49 ago)
  Input rate    : 0 bps (0 pps)
  Output rate   : 0 bps (0 pps)
  Active alarms : None
  Active defects: None
  Logical interface xe-0/1/0.0 (Index 88) (SNMP ifIndex 70)
    Flags: SNMP-Traps Encapsulation: ENET2
    Input packets : 0
    Output packets: 0
    Protocol eth-switch
    Flags: None

```

```

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
xe-0/1/0 brief**

```

user@switch> show interfaces 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
xe-4/1/0 detail**

```

user@switch> show interfaces 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-4/1/0 extensive**

```

user@switch> show interfaces xe-4/1/0 extensive
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 (00:14:01 ago)
Statistics last cleared: Never
Traffic statistics:
Input bytes :          5015472          36432 bps
Output bytes :           0          0 bps
Input packets:          3304          3 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: 3, 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 : None
Active defects : None
MAC statistics:
  Receive          Transmit
Total octets      5015472          0
Total packets    3304             0
Unicast packets  3304             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
Code violations   0

Packet Forwarding Engine configuration:
Destination slot: 4
Direction : Output
CoS transmit queue      Bandwidth      Buffer Priority
Limit
  %      bps      %      usec
  0 best-effort      95      9500000000      95      NA      low
none
  7 network-control  5       500000000       5       NA      low
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 lacp interfaces

Syntax	show lacp interfaces <i>interface-name</i>
Release Information	Command introduced in JUNOS 10.0 for EX Series switches.
Description	Display Link Aggregation Control Protocol (LACP) information about the specified aggregated Ethernet or Gigabit Ethernet interface.
Options	<p>none—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—<i>aex</i> ■ Gigabit Ethernet—<i>ge-fpc/pic/port</i>
Required Privilege Level	view
Related Topics	<ul style="list-style-type: none"> ■ Example: Configuring Aggregated Ethernet High-Speed Uplinks Between a Virtual Chassis Access Switch and a Virtual Chassis Distribution Switch on page 25 ■ Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between a Virtual Chassis Access Switch and a Virtual Chassis Distribution Switch on page 32 ■ Configuring Aggregated Ethernet Interfaces (CLI Procedure) on page 69 ■ Configuring Aggregated Ethernet LACP (CLI Procedure) on page 73 ■ Understanding Aggregated Ethernet Interfaces and LACP on page 7 ■ <i>JUNOS Software Network Interfaces Configuration Guide</i> at http://www.juniper.net/techpubs/software/junos/junos100/index.html
List of Sample Output	show lacp interfaces (Aggregated Ethernet) on page 169
Output Fields	Table 15 on page 166 lists the output fields for the show lacp interfaces command. Output fields are listed in the approximate order in which they appear.

Table 15: show lacp interfaces Output Fields

Field Name	Field Description
Aggregated interface	Aggregated Ethernet interface value.

Table 15: show lacp interfaces Output Fields (continued)

Field Name	Field Description
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 15: 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 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.

**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 distributing	CURRENT	Fast periodic	Collecting
ge-1/0/2 distributing	CURRENT	Fast periodic	Collecting
ge-1/0/1 (active) distributing	CURRENT	Fast periodic	Collecting
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 vrrp

Syntax show vrrp
 <brief | detail | extensive | summary>
 <interface *interface-name*>
 <track interfaces>

Release Information Statement introduced in JUNOS Release 10.0 for EX Series switches.

Description Display information and status about VRRP groups.

Options none—(Same as brief) Display brief status information about all VRRP interfaces.
 brief | detail | extensive | summary—(Optional) Display the specified level of output.
 interface *interface-name* —(Optional) Display information and status about the specified VRRP interface.
 track interfaces—(Optional) Display information and status about VRRP track interfaces.

Required Privilege Level view

- Related Topics**
- Configuring VRRP for IPv6 (CLI Procedure) on page 76
 - High Availability Features for EX Series Switches Overview on page 18

List of Sample Output

- show vrrp on page 174
- show vrrp brief on page 175
- show vrrp detail (IPv6) on page 175
- show vrrp detail (Route Track) on page 175
- show vrrp extensive on page 175
- show vrrp interface on page 177
- show vrrp summary on page 178
- show vrrp track detail on page 178
- show vrrp track summary on page 178

Output Fields Table 16 on page 170 lists the output fields for the `show vrrp` command. Output fields are listed in the approximate order in which they appear.

Table 16: show vrrp Output Fields

Field Name	Field Description	Level of Output
Interface	Name of the logical interface.	none, brief, extensive, summary
Interface index	Physical interface index number, which reflects its initialization sequence.	extensive
Groups	Total number of VRRP groups configured on the interface.	extensive

Table 16: show vrrp Output Fields (continued)

Field Name	Field Description	Level of Output
Active	Total number of VRRP groups that are active (that is, whose interface state is either up or down).	extensive
Interface VRRP PDU statistics	Nonerrored statistics for the logical interface: <ul style="list-style-type: none"> ■ Advertisement sent—Number of VRRP advertisement protocol data units (PDUs) that the interface has transmitted. ■ Advertisement received—Number of VRRP advertisement PDUs received by the interface. ■ Packets received—Number of VRRP packets received for VRRP groups on the interface. ■ No group match received—Number of VRRP packets received for VRRP groups that do not exist on the interface. 	extensive
Interface VRRP PDU error statistics	Errored statistics for the logical interface: <ul style="list-style-type: none"> ■ Invalid IPAH next type received—Number of packets received that use the IP Authentication Header protocol (IPAH) and that do not encapsulate VRRP packets. ■ Invalid VRRP ttl value received—Number of packets received whose IP time-to-live (TTL) value is not 255. ■ Invalid VRRP version received—Number of packets received whose VRRP version is not 2. ■ Invalid VRRP pdu type received—Number of packets received whose VRRP PDU type is not 1. ■ Invalid VRRP authentication type received—Number of packets received whose VRRP authentication is not none, simple, or md5. ■ Invalid VRRP IP count received—Number of packets received whose VRRP IP count exceeds 8. ■ Invalid VRRP checksum received—Number of packets received whose VRRP checksum does not match the calculated value. 	extensive
Physical interface	Name of the physical interface.	detail, extensive
Unit	Logical unit number.	All levels
Address	Address of the physical interface.	none, brief, detail, extensive
Index	Physical interface index number, which reflects its initialization sequence.	detail, extensive
SNMP ifIndex	SNMP index number for the physical interface.	detail, extensive
VRRP-Traps	Status of VRRP traps: Enabled or Disabled.	detail, extensive
Type and Address	Identifier for the address and the address itself: <ul style="list-style-type: none"> ■ lcl—Configured local interface address. ■ mas—Address of the master virtual router. This address is displayed only when the local interface is acting as a backup router. ■ vip—Configured virtual IP addresses. 	none, brief, summary

Table 16: show vrrp Output Fields (continued)

Field Name	Field Description	Level of Output
Interface state or Int state	State of the physical interface: <ul style="list-style-type: none"> ■ down—The device is present and the link is unavailable. ■ not present—The interface is configured, but no physical device is present. ■ unknown—The VRRP process has not had time to query the kernel about the state of the interface. ■ up—The device is present and the link is established. 	none, brief, extensive, summary
Group	VRRP group number.	none, brief, extensive, summary
State	VRRP state: <ul style="list-style-type: none"> ■ backup—The interface is acting as the backup router interface. ■ bringup—VRRP is just starting, and the physical device is not yet present. ■ idle—VRRP is configured on the interface and is disabled. This can occur when VRRP is first enabled on an interface whose link is established. ■ initializing—VRRP is initializing. ■ master—The interface is acting as the master router interface. ■ transition—The interface is changing between being the backup and being the master router. 	extensive
Priority	Configured VRRP priority for the interface.	detail, extensive
Advertisement interval	Configured VRRP advertisement interval.	detail, extensive
Authentication type	Configured VRRP authentication type: none, simple, or md5.	detail, extensive
Preempt	Whether preemption is allowed on the interface: yes or no.	detail, extensive
Accept-data mode	Whether the interface is configured to accept packets destined for the virtual IP address: yes or no.	detail, extensive
VIP count	Number of virtual IP addresses that have been configured on the interface.	detail, extensive
VIP	List of virtual IP addresses configured on the interface.	detail, extensive
Advertisement timer	Time until the advertisement timer expires.	detail, extensive
Master router	IP address of the interface that is acting as the master. If the VRRP interface is down, the output is N/A.	detail, extensive
Virtual router uptime	Time that the virtual router has been up.	detail, extensive
Master router uptime	Time that the master router has been up.	detail, extensive
Virtual MAC	MAC address associated with the virtual IP address.	detail, extensive
Tracking	Whether tracking is enabled or disabled.	detail, extensive

Table 16: show vrrp Output Fields (continued)

Field Name	Field Description	Level of Output
Current priority	Current operational priority for being the VRRP master.	detail, extensive
Configured priority	Configured base priority for being the VRRP master.	detail, extensive
Priority hold-time	Minimum time interval, in seconds, between successive changes to the current priority. Disabled indicates no minimum interval.	detail, extensive
Remaining-time	(track option only) Displays the time remaining in the priority hold-time interval.	detail
Interface tracking	Whether interface tracking is enabled or disabled. When enabled, the output also displays the number of tracked interfaces.	detail extensive
Interface/Tracked interface	Name of the tracked interface.	detail extensive
Int state/Interface state	Current operational state of the tracked interface: up or down .	detail, extensive
Int speed/Speed	Current operational speed, in bits per second, of the tracked interface.	detail, extensive
Incurred priority cost	Operational priority cost incurred due to the state and speed of this tracked interface. This cost is applied to the configured priority to obtain the current priority.	detail, extensive
Threshold	Speed below which the corresponding priority cost is incurred. In other words, when the speed of the interface drops below the threshold speed, the corresponding priority cost is incurred. An entry of down means that the corresponding priority cost is incurred when the interface is down.	detail, extensive
Route tracking	Whether route tracking is enabled or disabled. When enabled, the output also displays the number of tracked routes.	detail, extensive
Route count	The number of routes being tracked.	detail, extensive
Route	The IP address of the route being tracked.	detail, extensive
VRF name	The VPN routing and forwarding (VRF) routing instance that the tracked route is in.	detail, extensive
Route state	The state of the route being tracked: up , down , or unknown .	detail, extensive
Priority cost	Configured priority cost. This value is incurred when the interface speed drops below the corresponding threshold or when the tracked route goes down.	detail, extensive
Active	Whether the threshold is active (*). If the threshold is active, the corresponding priority cost is incurred.	detail, extensive
Group VRRP PDU statistics	Number of VRRP advertisements sent and received by the group.	extensive

Table 16: show vrrp Output Fields (continued)

Field Name	Field Description	Level of Output
Group VRRP PDU error statistics	<p>Errored statistics for the VRRP group:</p> <ul style="list-style-type: none"> ■ Bad authentication type received—Number of VRRP PDUs received with an invalid authentication type. The received authentication can be none, simple, or md5 and must be the same for all routers in the VRRP group. ■ Bad password received—Number of VRRP PDUs received with an invalid key (password). The password for simple authentication must be the same for all routers in the VRRP group ■ Bad MD5 digest received—Number of VRRP PDUs received for which the MD5 digest computed from the VRRP PDU differs from the digest expected by the VRRP instance configured on the router. ■ Bad advertisement timer received—Number of VRRP PDUs received with an advertisement time interval that is inconsistent with the one in use among the routers in the VRRP group. ■ Bad VIP count received—Number of VRRP PDUs whose virtual IP address counts differ from the count that has been configured on the VRRP instance. ■ Bad VIPADDR received—Number of VRRP PDUs whose virtual IP addresses differ from the list of virtual IP addresses configured on the VRRP instance. 	extensive
Group state transition statistics	<p>State transition statistics for the VRRP group:</p> <ul style="list-style-type: none"> ■ Idle to master transitions—Number of times that the VRRP instance transitioned from the idle state to the master state. ■ Idle to backup transitions—Number of times that the VRRP instance transitioned from the idle state to the backup state. ■ Backup to master transitions—Number of times that the VRRP instance transitioned from the backup state to the master state. ■ Master to backup transitions—Number of times that the VRRP instance transitioned from the master state to the backup state. 	extensive
VR state	<p>VRRP information:</p> <ul style="list-style-type: none"> ■ backup—The interface is acting as the backup router interface. ■ bringup—VRRP is just starting, and the physical device is not yet present. ■ idle—VRRP is configured on the interface and is disabled. This can occur when VRRP is first enabled on an interface whose link is established. ■ initializing—VRRP is initializing. ■ master—The interface is acting as the master router interface. ■ transition—The interface is changing between being the backup and being the master router. 	none, brief
Timer	<p>VRRP timer information:</p> <ul style="list-style-type: none"> ■ A—Time, in seconds, until the advertisement timer expires. ■ D—Time, in seconds, until the Master is Dead timer expires. 	none, brief

```

show vrrp user@host> show vrrp
Interface      State      Group  VR state  Timer  Type  Address
ge-0/0/0.121  up         1      master    A 1.052  1c1  gec0::12:1:1:1
    
```

```

vip ge80::12:1:1:99
vip gec0::12:1:1:99
ge-0/0/2.131 up 1 master A 0.364 1c1 gec0::13:1:1:1
vip ge80::13:1:1:99
vip gec0::13:1:1:99

```

show vrrp brief The output for the `show vrrp brief` command is identical to that for the `show vrrp` command. For sample output, see `show vrrp` on page 174.

show vrrp detail (IPv6) user@host> `show vrrp detail`
Physical interface: ge-0/0/0, Unit: 121, Vlan-id: 212, Address: gec0::12:1:1:1/120

```

Index: 67, SNMP ifIndex: 45, VRRP-Traps: enabled
Interface state: up, Group: 1, State: master
Priority: 200, Advertisement interval: 1, Authentication type: none
Preempt: yes, Accept-data mode: no, VIP count: 2, VIP: ge80::12:1:1:99,
gec0::12:1:1:99
Advertisement timer: 1.121s, Master router: ge80::12:1:1:1
Virtual router uptime: 00:03:47, Master router uptime: 00:03:41
Virtual MAC: 00:00:5e:00:02:01
Tracking: disabled

```

Physical interface: ge-0/0/2, Unit: 131, Vlan-id: 213, Address: gec0::13:1:1:1/120

```

Index: 69, SNMP ifIndex: 47, VRRP-Traps: enabled
Interface state: up, Group: 1, State: master
Priority: 200, Advertisement interval: 1, Authentication type: none
Preempt: yes, Accept-data mode: no, VIP count: 2, VIP: ge80::13:1:1:99,
gec0::13:1:1:99
Advertisement timer: 0.327s, Master router: ge80::13:1:1:1
Virtual router uptime: 00:03:47, Master router uptime: 00:03:41
Virtual MAC: 00:00:5e:00:02:01
Tracking: disabled

```

show vrrp detail (Route Track) user@host> `show vrrp detail`
Physical interface: ge-1/1/0, Unit: 0, Address: 30.30.30.30/24

```

Index: 67, SNMP ifIndex: 379, VRRP-Traps: enabled
Interface state: up, Group: 100, State: master
Priority: 150, Advertisement interval: 1, Authentication type: none
Preempt: yes, Accept-data mode: no, VIP count: 1, VIP: 30.30.30.100
Advertisement timer: 1.218s, Master router: 30.30.30.30
Virtual router uptime: 00:04:28, Master router uptime: 00:00:13
Virtual MAC: 00:00:5e:00:01:64
Tracking: enabled
Current priority: 150, Configured priority: 150
Priority hold-time: disabled
Interface tracking: disabled
Route tracking: enabled, Route count: 1
Route          VRF name      Route state  Priority cost
192.168.40.0/22 default       up           30

```

show vrrp extensive user@host> `show vrrp extensive`
Interface: ge-0/0/0.121, Interface index: 67, Groups: 1, Active : 1

```

Interface VRRP PDU statistics
Advertisement sent           :           188
Advertisement received      :              0

```

```

Packets received                :           0
No group match received         :           0
Interface VRRP PDU error statistics
Invalid IPAH next type received :           0
Invalid VRRP TTL value received :           0
Invalid VRRP version received  :           0
Invalid VRRP PDU type received :           0
Invalid VRRP authentication type received:           0
Invalid VRRP IP count received :           0
Invalid VRRP checksum received :           0

Physical interface: ge-0/0/0, Unit: 121, Vlan-id: 212, Address: gec0::12:1:1:1/120

Index: 67, SNMP ifIndex: 45, VRRP-Traps: enabled
Interface state: up, Group: 1, State: master
Priority: 200, Advertisement interval: 1, Authentication type: none
Preempt: yes, Accept-data mode: no, VIP count: 2, VIP: ge80::12:1:1:99,
gec0::12:1:1:99
Advertisement timer: 1.034s, Master router: ge80::12:1:1:1
Virtual router uptime: 00:04:04, Master router uptime: 00:03:58
Virtual MAC: 00:00:5e:00:02:01
Tracking: disabled
Group VRRP PDU statistics
  Advertisement sent           :          188
  Advertisement received       :           0
Group VRRP PDU error statistics
Bad authentication type received:           0
Bad password received          :           0
Bad MD5 digest received        :           0
Bad advertisement timer received:           0
Bad VIP count received         :           0
Bad VIPADDR received           :           0
Group state transition statistics
  Idle to master transitions    :           0
  Idle to backup transitions    :           1
  Backup to master transitions  :           1
  Master to backup transitions  :           0

Interface: ge-0/0/2.131, Interface index: 69, Groups: 1, Active : 1
Interface VRRP PDU statistics
  Advertisement sent           :          186
  Advertisement received       :           0
  Packets received             :           0
  No group match received      :           0
Interface VRRP PDU error statistics
Invalid IPAH next type received :           0
Invalid VRRP TTL value received :           0
Invalid VRRP version received  :           0
Invalid VRRP PDU type received :           0
Invalid VRRP authentication type received:           0
Invalid VRRP IP count received :           0
Invalid VRRP checksum received :           0

Physical interface: ge-0/0/2, Unit: 131, Vlan-id: 213, Address: gec0::13:1:1:1/120

Index: 69, SNMP ifIndex: 47, VRRP-Traps: enabled
Interface state: up, Group: 1, State: master
Priority: 200, Advertisement interval: 1, Authentication type: none
Preempt: yes, Accept-data mode: no, VIP count: 2, VIP: ge80::13:1:1:99,
gec0::13:1:1:99
Advertisement timer: 0.396s, Master router: ge80::13:1:1:1

```

```

Virtual router uptime: 00:04:04, Master router uptime: 00:03:58
Virtual MAC: 00:00:5e:00:02:01
Tracking: disabled
Group VRRP PDU statistics
  Advertisement sent      :      186
  Advertisement received  :         0
Group VRRP PDU error statistics
  Bad authentication type received:    0
  Bad password received      :         0
  Bad MD5 digest received     :         0
  Bad advertisement timer received:    0
  Bad VIP count received      :         0
  Bad VIPADDR received       :         0
Group state transition statistics
  Idle to master transitions :         0
  Idle to backup transitions :         1
  Backup to master transitions :        1
  Master to backup transitions :         0

```

show vrrp interface

```

user@host> show vrrp interface
Interface: ge-0/0/0.121, Interface index: 67, Groups: 1, Active : 1
Interface VRRP PDU statistics
  Advertisement sent      :      205
  Advertisement received  :         0
  Packets received        :         0
  No group match received :         0
Interface VRRP PDU error statistics
  Invalid IPAH next type received :    0
  Invalid VRRP TTL value received :    0
  Invalid VRRP version received   :    0
  Invalid VRRP PDU type received  :    0
  Invalid VRRP authentication type received: 0
  Invalid VRRP IP count received  :    0
  Invalid VRRP checksum received  :    0

Physical interface: ge-0/0/0, Unit: 121, Vlan-id: 212, Address: gec0::12:1:1:1/120

Index: 67, SNMP ifIndex: 45, VRRP-Traps: enabled
Interface state: up, Group: 1, State: master
Priority: 200, Advertisement interval: 1, Authentication type: none
Preempt: yes, Accept-data mode: no, VIP count: 2, VIP: ge80::12:1:1:99,
gec0::12:1:1:99
Advertisement timer: 0.789s, Master router: ge80::12:1:1:1
Virtual router uptime: 00:04:26, Master router uptime: 00:04:20
Virtual MAC: 00:00:5e:00:02:01
Tracking: disabled
Group VRRP PDU statistics
  Advertisement sent      :      205
  Advertisement received  :         0
Group VRRP PDU error statistics
  Bad authentication type received:    0
  Bad password received      :         0
  Bad MD5 digest received     :         0
  Bad advertisement timer received:    0
  Bad VIP count received      :         0
  Bad VIPADDR received       :         0
Group state transition statistics
  Idle to master transitions :         0
  Idle to backup transitions :         1
  Backup to master transitions :         1
  Master to backup transitions :         0

```

```

show vrrp summary user@host> show vrrp summary
Interface      State      Group  VR state   Type   Address
ge-4/1/0.0    up         1      backup    lc1    10.57.0.2
vip            10.57.0.100
    
```

```

show vrrp track detail user@host> show vrrp track detail
Tracked interface: ae1.211
State: up, Speed: 400m
Incurred priority cost: 0
Threshold      Priority cost  Active
400m           10
300m           60
200m           110
100m           160
down           190
Tracking VRRP interface: ae0.210, Group: 1
VR State: master
Current priority: 200, Configured priority: 200
Priority hold-time: disabled, Remaining-time: 50.351
    
```

```

show vrrp track summary user@host> show vrrp track summary
Track if      State   Speed   VRRP if  Group  VR State  Current priority
ae1.211       up     400m   ae0.210  1      master    200
    
```

Part 2

Power over Ethernet

- Power over Ethernet (PoE)—Overview on page 181
- Examples: PoE Configuration on page 185
- Configuring PoE on page 193
- Verifying PoE Configuration on page 197
- Configuration Statements for PoE on page 199
- Operational Mode Commands for PoE on page 209

Chapter 8

Power over Ethernet (PoE)—Overview

- PoE and EX Series Switches Overview on page 181

PoE and EX Series Switches Overview

Power over Ethernet (PoE) is the implementation of IEEE 802.3af, allowing both data and electric power to pass over a copper Ethernet LAN cable. This technology allows 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.

This topic covers:

- PoE and Power Supply Units in EX Series Switches on page 181
- Power Management Mode on page 182
- Classes of Powered Devices on page 182
- Global and Specific PoE Parameters on page 183

PoE and Power Supply Units in EX Series Switches

Juniper Networks EX Series Ethernet Switch models provide either 8, 24, or 48 PoE ports. The total number of PoE ports for an EX Series switch can be extended by inserting additional PoE cards.

Power supply units with three different power capacities are available for use with the EX Series switches:

- 320-W power supply unit: Supports 8 ports of PoE power at 15.4 W per port, plus system power.
- 600-W power supply unit: Supports 24 ports of PoE power at 15.4 W per port, plus system power.
- 930-W power supply unit: Supports 48 ports of PoE power at 15.4 W per port, plus system power.



NOTE: PoE is not supported on a switch using DC power.

All 802.3af-compliant powered devices require no more than 12.95 watts. Thus, if you follow the recommended guidelines for selecting power supply units to support the number of PoE ports, the switch should be able to supply power to all connected powered devices. If you install a higher capacity power supply unit on a switch model that has only 8 PoE ports, it does not extend PoE capabilities to the non-PoE ports.

Power Management Mode

You can use the power management mode to determine the number of interfaces that can be provided with power. The following two factors constitute the power management mode:

- Per port limit (PPL)—The factor that decides the maximum power consumption permitted on a particular interface. If the power consumption by the powered device exceeds the specified value, PoE is shut down over that interface.
- Power allocated for each interface—The factor that ensures that a certain amount of power is reserved for an individual interface from the total power budget for all interfaces. If at any point the total of the allocated power for all interfaces exceeds the total budget, the lower priority interfaces are turned off and the power allocated for those interfaces drops to 0.

There are two modes of power management:

- Static—In this mode the power allocated for each interface can be configured. The PPL value is the maximum value configured per interface.
- Class—In this mode the power allocation for interfaces is determined based on the class of powered device connected. PPL is the maximum power value of the class of the powered device connected to the interface. The power allocated per interface is the maximum power of the powered device class, except for classes 0 and 3. For class 0 and class 3 powered devices, the momentary power consumption is considered as the power allocated for that interface. Therefore, PPL and power allocated per interface values change based on the powered device connected to the interface.

Classes of Powered Devices

A powered device is classified based on the maximum power that it draws across all input voltages and operational modes. The most common class is 0, in which the switch allows a maximum draw of 15.4 W per port. The switch provides 15.4 W at the port in order to guarantee enough power to run a device, after accounting for line loss. For example, $15.4 \text{ W} - \text{power loss (16\%)} = 12.95 \text{ W}$. Table 17 on page 182 lists the classes of powered devices and associated power levels.

Table 17: Class of Powered Device and Power Levels

Class	Usage	Minimum Power Levels Output from PoE Port	Range of Maximum Power required by the Powered Device
0	Default	15.4 W	0.44 through 12.95 W
1	Optional	4.0 W	0.44 through 3.84 W

Table 17: Class of Powered Device and Power Levels (continued)

Class	Usage	Minimum Power Levels Output from PoE Port	Range of Maximum Power required by the Powered Device
2	Optional	7.0 W	6.49 through 12.95 W
3	Optional	15.4 W	6.49 through 12.95 W

Global and Specific PoE Parameters

All EX Series switches with PoE ports have a PoE controller. The PoE controller keeps track of the switch's power consumption and distributes the available power to individual PoE ports. You can set the PoE controller to reserve a limited amount of power (up to 19 W) to handle a power spike. The default is that no power is kept on reserve.

The factory default configuration creates a PoE interface for all the PoE ports on the switch. You can specify maximum power, priority, and telemetries for each PoE interface.

- **maximum-power**—This setting defaults to 15.4 W. If you follow the recommended guidelines for the installed power supply unit (see Table 17 on page 182), the switch should be able to provide sufficient power for all PoE ports using the default power setting.
- **priority**—This setting defaults to **low**. If a port is set as high priority and a situation arises where there is not sufficient power for all the PoE ports, the available power is directed to the higher priority port(s). If the switch needs to shut down powered devices because a power supply fails and there is insufficient power, low priority devices are shut before high priority powered devices. Thus, security cameras, emergency phones, and other high priority phones should be set to high priority.
- **telemetries**—This setting allows you to monitor per port PoE power consumption. It is not included in the default PoE configuration.

Related Topics

- EX Series Switches Interfaces Overview on page 3
- Example: Configuring PoE Interfaces on an EX Series Switch on page 185
- Example: Configuring PoE Interfaces with Different Priorities on an EX Series Switch on page 188

Chapter 9

Examples: PoE Configuration

- Example: Configuring PoE Interfaces on an EX Series Switch on page 185
- Example: Configuring PoE Interfaces with Different Priorities on an EX Series Switch on page 188

Example: Configuring PoE Interfaces on an EX Series Switch

All EX Series switches except the EX4200-24F model provide Power over Ethernet (PoE) ports. The PoE ports supply electric power over the same ports that are used to connect network devices and allow you to plug in devices that require both network connectivity and electric power, such as VoIP phones, wireless access points, and some IP cameras. The factory default configuration specifies PoE interfaces for the PoE ports. Therefore, you do not need to configure PoE unless you wish to modify the default values or disable a specific PoE interface.

This example describes a default configuration of PoE interfaces on an EX Series switch:

- Requirements on page 185
- Overview and Topology on page 186
- Configuration on page 186
- Verification on page 187
- Troubleshooting on page 187

Requirements

This example uses the following software and hardware components:

- JUNOS Release 9.0 or later for EX Series switches
- One EX4200 switch

Before you configure PoE, be sure you have:

- Performed the initial switch configuration. See *Connecting and Configuring an EX Series Switch (CLI Procedure)* or *Connecting and Configuring an EX Series Switch (J-Web Procedure)* for details.

Overview and Topology

The topology used in this example consists of one EX4200-24T switch, which has a total of 24 ports. Eight of the ports support PoE, which means they provide both network connectivity and electric power for devices such as VoIP phones, wireless access points, and some IP security cameras. The remaining 16 ports provide only network connectivity. You use the standard ports to connect devices that have their own power sources, such as desktop and laptop computers, printers, and servers. Table 18 on page 186 details the topology used in this configuration example.

Table 18: Components of the PoE Configuration Topology

Property	Settings
Switch hardware	EX4200-E-24T switch, with 24 Gigabit Ethernet ports: 8 PoE ports (ge-0/0/0 through ge-0/0/7) and 16 non-PoE ports (ge-0/0/8 through ge-0/0/23)
VLAN name	default
Connection to a wireless access point (requires PoE)	ge-0/0/0
Connections to Avaya IP telephone—with integrated hub, to connect phone and desktop PC to a single port (requires PoE)	ge-0/0/1 through ge-0/0/7
Direct connections to desktop PCs, file servers, integrated printer/fax/copier machines (no PoE required)	ge-0/0/8 through ge-0/0/20
Unused ports (for future expansion)	ge-0/0/21 through ge-0/0/23

Configuration

To enable the default PoE configuration on the switch:

CLI Quick Configuration By default, PoE interfaces are created for all PoE ports and PoE is enabled. You can simply connect powered devices to the PoE ports.

Step-by-Step Procedure To use the PoE interfaces with default values:

1. Make sure the switch is powered on.
2. Connect the wireless access point to switch port ge-0/0/0.
3. Connect the eight Avaya phones to switch ports ge-0/0/1 through ge-0/0/7.

Verification

To verify that PoE interfaces have been created and are operational, perform this task:

- Verifying That the PoE Interfaces Have Been Created on page 187

Verifying That the PoE Interfaces Have Been Created

Purpose Verify that the PoE interfaces have been created on the switch.

Action List all the PoE interfaces configured on the switch:

```
user@switch>
                               show poe interface

Interface  Enabled  status  max-power  priority  power-consumption  Class
ge-0/0/0   Enabled  ON      15.4W     Low       12.95W             0
ge-0/0/1   Enabled  ON      15.4W     Low       12.95W             0
ge-0/0/2   Enabled  ON      15.4W     Low       12.95W             0
ge-0/0/3   Enabled  ON      15.4W     Low       12.95W             0
ge-0/0/4   Enabled  ON      15.4W     Low       12.95W             0
ge-0/0/5   Enabled  ON      15.4W     Low       12.95W             0
ge-0/0/6   Enabled  ON      15.4W     Low       12.95W             0
ge-0/0/7   Enabled  ON      15.4W     Low       12.95W             0
```

Meaning The `show poe interface` command lists PoE interfaces configured on the switch, with their status, priority, power consumption, and class. This output shows that eight interfaces have been created with default values and are consuming power at the expected rates.

Troubleshooting

Troubleshooting PoE Interfaces

Problem The PoE port is not supplying power to the port.

Solution Check for the following:

Items to Check	Explanation
Is the switch a full PoE model or partial PoE?	If you are using a partial PoE model, only interfaces <code>ge-0/0/0</code> through <code>ge-0/0/7</code> can function as PoE ports.
Has the PoE interface been disabled for that port?	Use the <code>show poe interface</code> command to check PoE interface status.
Is the cable properly seated in the port socket?	Check the hardware.
Enable <code>telemetries</code> for the interface.	Check the history of power consumption on the interface by using the <code>show poe telemetries interface</code> command.

- Related Topics**
- Example: Configuring PoE Interfaces with Different Priorities on an EX Series Switch on page 188
 - Configuring PoE (CLI Procedure) on page 193
 - Configuring PoE (J-Web Procedure) on page 195

Example: Configuring PoE Interfaces with Different Priorities on an EX Series Switch

EX Series switches provide Power over Ethernet (PoE) ports, which supply electric power over the same ports that are used to connect network devices. These ports allow you to plug in devices that need both network connectivity and electric power, such as VoIP phones, wireless access points, and some IP cameras. You can configure a particular PoE interface to have a high priority setting. If a port is set as high priority and a situation arises where there is not sufficient power for all the PoE ports, the available power is directed to the higher priority ports. If the switch needs to shut down powered devices because a power supply fails and there is insufficient power, low priority devices are shut down before high priority powered devices. Thus, security cameras, emergency phones, and other high priority phones should be set to high priority.

This example describes how to configure a few high priority PoE interfaces for an EX Series switch (by default, interfaces are set to low priority):

- Requirements on page 188
- Overview and Topology on page 188
- Configuration on page 189
- Verification on page 190
- Troubleshooting on page 191

Requirements

This example uses the following software and hardware components:

- JUNOS Release 9.0 or later for EX Series switches
- One EX4200 switch

Before you configure PoE, be sure you have:

- Performed the initial switch configuration. See *Connecting and Configuring an EX Series Switch (CLI Procedure)* or *Connecting and Configuring an EX Series Switch (J-Web Procedure)* for details.

Overview and Topology

The topology used in this example consists of one EX4200-24T switch, which has a total of 24 ports. Eight of the ports support PoE, which means they provide both network connectivity and electric power for devices such as VoIP telephones, wireless

access points, and some IP security cameras. The remaining 16 ports provide only network connectivity. You use the standard ports to connect devices that have their own power sources, such as desktop and laptop computers, printers, and servers. Table 19 on page 189 details the topology used in this configuration example.

Table 19: Components of the PoE Configuration Topology

Property	Settings
Switch hardware	EX4200-E-24T switch, with 24 Gigabit Ethernet ports: 8 PoE ports (ge-0/0/0 through ge-0/0/7) and 16 non-PoE ports (ge-0/0/8 through ge-0/0/23)
VLAN name	default
Connection to a wireless access point (requires PoE)	ge-0/0/0
Security IP Cameras (require PoE)	ge-0/0/1 and ge-0/0/2 high
Emergency VoIP phone (requires PoE)	ge-0/0/3 high
VoIP phone in Executive Office (requires PoE)	ge-0/0/4 high
Other VoIP phones (require PoE)	ge-0/0/5 through ge-0/0/7
Direct connections to desktop PCs, file servers, integrated printer/fax/copier machines (no PoE required)	ge-0/0/8 through ge-0/0/20
Unused ports (for future expansion)	ge-0/0/21 through ge-0/0/23

Configuration

Configure Power over Ethernet Interfaces:

CLI Quick Configuration By default, PoE interfaces are created for all PoE ports and PoE is enabled. The default priority for PoE interfaces is low.

To quickly configure PoE with some interfaces set to high priority and others to the default low priority, and to include a description of the interfaces, copy the following commands and paste them into the switch terminal window:

```
[edit]
set poe interface ge-0/0/1 priority high telemetries
set poe interface ge-0/0/2 priority high telemetries
set poe interface ge-0/0/3 priority high telemetries
set poe interface ge-0/0/4 priority high telemetries
set poe interface all
set interfaces ge-0/0/0 description "wireless access point"
set interfaces ge-0/0/1 description "security camera front door"
set interfaces ge-0/0/2 description "security camera back door"
set interfaces ge-0/0/3 description "emergency phone"
set interfaces ge-0/0/4 description "Executive Office VoIP phone"
set interfaces ge-0/0/5 description "staff VoIP phone"
set interfaces ge-0/0/6 description "staff VoIP phone"
set interfaces ge-0/0/7 description "staff VoIP phone"
```

Step-by-Step Procedure To configure PoE interfaces with different priorities:

1. Configure the PoE interfaces at the [edit poe] hierarchy level with some interfaces set to high priority and others to the default low priority, thus enabling the logging of per-port power consumption for the high priority ports.

```
[edit poe]
user@switch# set interface ge-0/0/1 priority high telemetries
user@switch# set interface ge-0/0/2 priority high telemetries
user@switch# set interface ge-0/0/3 priority high telemetries
user@switch# set interface ge-0/0/4 priority high telemetries
user@switch# set interface all
```

2. Specify a description for the PoE interfaces:

```
[edit interfaces]
user@switch# set ge-0/0/0 description "wireless access point"
user@switch# set ge-0/0/1 description "security camera front door"
user@switch# set ge-0/0/2 description "security camera back door"
user@switch# set ge-0/0/3 description "emergency phone"
user@switch# set ge-0/0/4 description "Executive Office VoIP phone"
user@switch# set ge-0/0/5 description "staff VoIP phone"
user@switch# set ge-0/0/6 description "staff VoIP phone"
user@switch# set ge-0/0/7 description "staff VoIP phone"
```

3. Connect the wireless access point to switch interface `ge-0/0/0`. This interface is PoE-enabled for the default settings based on the factory configuration. Telemetries are not enabled.
4. Connect the two security cameras to switch interfaces `ge-0/0/1` and `ge-0/0/2`. These interfaces are set to high priority with telemetries enabled.
5. Connect the emergency VoIP phone to switch interface `ge-0/0/3`. This interface is set to high priority with telemetries enabled.
6. Connect the Executive Office VoIP phone to switch interface `ge-0/0/4`. This interface is set to high priority with telemetries enabled.

Results Connect the staff VoIP phones to switch interfaces `ge-0/0/5` through `ge-0/0/7`. These interfaces are set to the default values. Telemetries are not enabled.

Verification

To verify that PoE interfaces have been created and are operational, perform the following tasks:

- Verifying That the PoE Interfaces Have Been Created with Desired Priorities on page 190

Verifying That the PoE Interfaces Have Been Created with Desired Priorities

Purpose Verify that the PoE interfaces on the switch are now set to the desired priority settings.

Action List all the PoE interfaces configured on the switch:

```
user@switch>
                               show poe interface

Interface  Enabled  Status  Max-Power  Priority  Power-Consumption  Class
ge-0/0/0   Enabled  ON      15.4W      Low      12.95W             0
ge-0/0/1   Enabled  ON      15.4W      High     12.95W             0
ge-0/0/2   Enabled  ON      15.4W      High     12.95W             0
ge-0/0/3   Enabled  ON      15.4W      High     12.95W             0
ge-0/0/4   Enabled  ON      15.4W      Low      12.95W             0
ge-0/0/5   Enabled  ON      15.4W      Low      12.95W             0
ge-0/0/6   Enabled  ON      15.4W      Low      12.95W             0
ge-0/0/7   Enabled  OFF     15.4W      Low      0 W                 0
```

Meaning The `show poe interface` command lists PoE interfaces configured on the switch, with their status, priority, power consumption, and class. This output shows that eight PoE interfaces are enabled. Interfaces `ge-0/0/1` through `ge-0/0/3` are configured as priority high. The remaining interfaces are configured with the default values.

Troubleshooting

Troubleshooting PoE Interfaces

Problem The PoE port is not supplying power to the port.

Solution Check for the following:

Items to Check	Explanation
Is the switch a full PoE model or partial PoE?	If you are using a partial PoE model, only interfaces <code>ge-0/0/0</code> through <code>ge-0/0/7</code> can function as PoE ports.
Has the PoE interface been disabled for that port?	Use the <code>show poe interface</code> command to check PoE interface status.
Is the cable properly seated in the port socket?	Check the hardware.
Enable <code>telemetries</code> for the interface.	Check the history of power consumption on the interface by using the <code>show poe telemetries interface</code> command.

- Related Topics**
- Example: Configuring PoE Interfaces on an EX Series Switch on page 185
 - Configuring PoE (CLI Procedure) on page 193
 - Configuring PoE (J-Web Procedure) on page 195

Chapter 10

Configuring PoE

- Configuring PoE (CLI Procedure) on page 193
- Configuring PoE (J-Web Procedure) on page 195

Configuring PoE (CLI Procedure)

EX Series switch models provide either 8, 24, or 48 PoE ports, which supply electric power over the same ports that are used to connect network devices. These ports allow you to plug in devices that require both network connectivity and electric power, such as VoIP phones, wireless access points, and some IP cameras.

The factory default configuration for EX Series switches specifies and enables PoE interfaces for the PoE ports.

To configure PoE using the CLI:

1. Enable PoE:
 - For all PoE interfaces:

```
[edit]  
user@switch# set poe interface all
```

- For a specific PoE interface:

```
[edit]  
user@switch# set poe interface ge-0/0/0
```

2. By default the power management mode is **static**. To change the power management mode to **class**:

```
[edit]  
user@switch# set poe management class
```



NOTE: When the power management mode is set to class, the maximum power value is overridden by the maximum power value of the class of power device connected.

3. Set the power priority:

- For all PoE interfaces:

```
[edit]
user@switch# set poe interface all priority low
```

- For a specific PoE interface:

```
[edit]
user@switch# set poe interface ge-0/0/0 priority high
```

4. Set the maximum PoE wattage available (the default is 15.4):

- For all PoE interfaces:

```
[edit]
user@switch# set poe interface all maximum-power 14
```

- For a specific PoE interface:

```
[edit]
user@switch# set poe interface ge-0/0/0 maximum-power 12.8
```

5. Enable logging of PoE power consumption with the default **telemetries** settings:

- For all PoE interfaces:

```
[edit]
user@switch# set poe interface all telemetries
```

- For a specific PoE interface:

```
[edit]
user@switch# set poe interface ge-0/0/0 telemetries
```

6. Reserve a specified wattage of power for the switch in case of a spike in PoE consumption (the default is 0):

```
[edit]
user@switch# set poe guard-band 15
```

- Related Topics**
- Configuring PoE (J-Web Procedure) on page 195
 - Example: Configuring PoE Interfaces on an EX Series Switch on page 185
 - Example: Configuring PoE Interfaces with Different Priorities on an EX Series Switch on page 188
 - Monitoring PoE on page 197
 - PoE and EX Series Switches Overview on page 181

Configuring PoE (J-Web Procedure)

EX Series switch models provide either 8, 24, or 48 PoE ports, which supply electric power over the same ports that are used to connect network devices. These ports allow you to plug in devices that require both network connectivity and electric power, such as VoIP phones, wireless access points, and some IP cameras. Using the Power over Ethernet (PoE) configuration page, you can modify the settings of all interfaces that are PoE-enabled.

To modify PoE settings:

1. In the **Configure** menu, select **Power over Ethernet**.

The page displays a list of all interfaces except uplink ports. Specific operational details about an interface are displayed in the **Details** section of the page. The details include the PoE Operational Status and Port class.

2. Click one:
 - **Edit** — Changes PoE settings for the selected port as described in Table 20 on page 195.
 - **System Settings** — Modifies general PoE settings as described in Table 21 on page 195.

Table 20: PoE Edit Settings

Field	Description	Your Action
Enable PoE	Specifies that PoE is enabled on the interface.	Select this option to enable PoE on the interface.
Priority	Lists the power priority (Low or High) configured on ports enabled for PoE.	Set the priority as High or Low .
Maximum Power	Specifies the maximum PoE wattage available to provision active PoE ports on the switch.	Select a value in watts. If no value is specified, the default is 15.4.

Table 21: System Settings

Field	Description	Your Action
PoE Management	<p>Specifies the power management mode. The options are: static and class.</p> <p>NOTE: When the power management mode is set to class, the maximum power value is overridden by the maximum power value of the class of power device that is connected to the switch on the PoE port.</p>	By default the power management mode is static . Select class to change the power management mode.
Guard Band (watts)	Specifies the band to control power availability on the switch.	Enter a value to set the guard band value in watts. The default value is 0.

- Related Topics**
- Configuring PoE (CLI Procedure) on page 193
 - Example: Configuring PoE Interfaces on an EX Series Switch on page 185
 - Example: Configuring PoE Interfaces with Different Priorities on an EX Series Switch on page 188
 - Monitoring PoE on page 197
 - PoE and EX Series Switches Overview on page 181

Chapter 11

Verifying PoE Configuration

- Monitoring PoE on page 197
- Verifying Status of PoE Interfaces on an EX Series Switch on page 198

Monitoring PoE

Purpose Use the monitoring functionality to view real-time data of the power consumed by each PoE interface, and to enable and configure Telemetries values. When Telemetries is enabled, the software measures the power consumed by each interface and stores the data for future reference.

Action To monitor PoE using the J-Web interface, select **Monitor > Power over Ethernet**.

To monitor PoE using the CLI:

- To display the real-time PoE status for all PoE interfaces, enter `show poe interface .`
- To display the real-time PoE status for a specific PoE interface, enter `show poe interface interface-name .`

The `show poe interface` command displays the power consumption of the interface at the moment that the command is issued.

To monitor the PoE interface's power consumption over a period of time, you can enable telemetries for the interface with the `telemetries` configuration statement. When Telemetries is enabled, you can display the log of the interface's power consumption by using the CLI command:

```
show poe telemetries interface interface-name all| x
```

Meaning In the J-Web interface the PoE Monitoring screen is divided into two parts. The top half of the screen displays real-time data of the power consumed by each interface and a list of ports that utilize maximum power.

Select a particular interface to view a graph of the power consumed by the selected interface.

The bottom half of the screen displays telemetries values for interfaces. The telemetry status displays whether telemetry has been enabled on the interface. Click the **Show Graph** button to view a graph of the telemetries. The graph can be based on power or voltage. To modify telemetries values, click **Edit**. Specify Interval in minutes,

Duration in hours, and select **Log Telemetries** to enable telemetries on the selected interface.

- Related Topics**
- Configuring PoE (CLI Procedure) on page 193
 - Configuring PoE (J-Web Procedure) on page 195
 - Example: Configuring PoE Interfaces on an EX Series Switch on page 185
 - Example: Configuring PoE Interfaces with Different Priorities on an EX Series Switch on page 188
 - Verifying Status of PoE Interfaces on an EX Series Switch on page 198

Verifying Status of PoE Interfaces on an EX Series Switch

Purpose Verify that the PoE interfaces on the switch are enabled and set to the desired priority settings.

Action List all the PoE interfaces configured on the switch:

```
user@switch> show poe interface
Interface  Enabled  Status  Max-Power  Priority  Power-Consumption  Class
ge-0/0/0   Enabled  ON      15.4W      Low      12.95W              0
ge-0/0/1   Enabled  ON      15.4W      High     12.95W              0
ge-0/0/2   Enabled  ON      15.4W      High     12.95W              0
ge-0/0/3   Enabled  ON      15.4W      High     12.95W              0
ge-0/0/4   Enabled  ON      15.4W      Low      12.95W              0
ge-0/0/5   Enabled  ON      15.4W      Low      12.95W              0
ge-0/0/6   Enabled  ON      15.4W      Low      12.95W              0
ge-0/0/7   Enabled  OFF     15.4W      Low      0 W                 0
```

Meaning The `show poe interface` command lists PoE interfaces configured on the switch, with their status, priority, power consumption, and class. This command has been executed on a switch with partial PoE (8 PoE ports). The output shows that all eight PoE interfaces are enabled. Interfaces `ge-0/0/1` through `ge-0/0/3` are configured as **priority high**. The remaining interfaces were configured with the default values.

- Related Topics**
- Configuring PoE (CLI Procedure) on page 193
 - Configuring PoE (J-Web Procedure) on page 195
 - Example: Configuring PoE Interfaces on an EX Series Switch on page 185
 - Example: Configuring PoE Interfaces with Different Priorities on an EX Series Switch on page 188
 - Monitoring PoE on page 197

Chapter 12

Configuration Statements for PoE

- [edit poe] Configuration Statement Hierarchy on page 199

[edit poe] Configuration Statement Hierarchy

```
poe {
  guard-band watts;
  interface (all | interface-name) {
    disable;
    maximum-power watts;
    priority value;
    telemetries {
      disable;
      duration hours;
      interval minutes;
    }
  }
  management type;
}
```

- Related Topics**
- Example: Configuring PoE Interfaces on an EX Series Switch on page 185
 - Example: Configuring PoE Interfaces with Different Priorities on an EX Series Switch on page 188
 - Configuring PoE (CLI Procedure) on page 193
 - Configuring PoE (J-Web Procedure) on page 195
 - PoE and EX Series Switches Overview on page 181

disable

Syntax	disable;
Hierarchy Level	[edit poe interface (all <i>interface-name</i>)], [edit poe interface (all <i>interface-name</i>) telemetries]
Release Information	Statement introduced in JUNOS Release 9.0 for EX Series switches.
Description	<p>Disables the PoE capabilities of this port. The port operates as a standard network access port. If the disable statement is specified after the telemetries statement, it disables the logging of PoE power consumption for this port.</p> <p>To disable the monitoring and retain the stored configuration values for interval and duration for possible future use, you can specify the disable substatement in the stanza for telemetries.</p>
Default	The PoE capabilities are automatically enabled when a PoE interface is set. If the telemetries statement is specified, monitoring of PoE per-port power consumption is enabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Topics	<ul style="list-style-type: none"> ■ Example: Configuring PoE Interfaces on an EX Series Switch on page 185 ■ Example: Configuring PoE Interfaces with Different Priorities on an EX Series Switch on page 188 ■ Configuring PoE (CLI Procedure) on page 193 ■ Configuring PoE (J-Web Procedure) on page 195 ■ PoE and EX Series Switches Overview on page 181

duration

Syntax	<code>duration hours;</code>
Hierarchy Level	[edit poe interface (all <i>interface-name</i>) telemetries]
Release Information	Statement introduced in JUNOS Release 9.0 for EX Series switches.
Description	Modify the duration for logging telemetries if you are monitoring the per-port power consumption for PoE interfaces.
Options	<p><i>hours</i> —Hours the logging continues.</p> <p>Range: 1 through 24 hours</p> <p>Default: 1 hour</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Topics	<ul style="list-style-type: none"> ■ Example: Configuring PoE Interfaces on an EX Series Switch on page 185 ■ Example: Configuring PoE Interfaces with Different Priorities on an EX Series Switch on page 188 ■ Configuring PoE (CLI Procedure) on page 193 ■ Configuring PoE (J-Web Procedure) on page 195 ■ PoE and EX Series Switches Overview on page 181

guard-band

Syntax	guard-band <i>watts</i> ;
Hierarchy Level	[edit poe]
Release Information	Statement introduced in JUNOS Release 9.0 for EX Series switches.
Description	Reserve the specified amount of power for the switch in case of a spike in PoE consumption.
Default	0 W
Options	<i>watts</i> —Amount of power to be reserved for the switch in case of a spike in PoE consumption. Range: 0 through 19 W Default: 0 W
Required Privilege Level	router—To view this statement in the configuration. router-control—To add this statement to the configuration.
Related Topics	<ul style="list-style-type: none">■ Example: Configuring PoE Interfaces on an EX Series Switch on page 185■ Example: Configuring PoE Interfaces with Different Priorities on an EX Series Switch on page 188■ Configuring PoE (CLI Procedure) on page 193■ Configuring PoE (J-Web Procedure) on page 195■ PoE and EX Series Switches Overview on page 181

interface

Syntax interface (all | *interface-name*) {
 disable;
 maximum-power *watts*;
 priority *value*;
 telemetries {
 disable;
 interval *minutes*;
 duration *hours*;
 }
 }

Hierarchy Level [edit poe]

Release Information Statement introduced in JUNOS Release 9.0 for EX Series switches.

Description Enable a PoE interface for a PoE port. An interface must be enabled in order for the port to provide power to a connected powered device.

Default The PoE interface is enabled by default.

Options all—All interfaces on the switch.

interface-name —Name of the specific interface.

The remaining statements are explained separately.

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

- Related Topics**
- Example: Configuring PoE Interfaces on an EX Series Switch on page 185
 - Example: Configuring PoE Interfaces with Different Priorities on an EX Series Switch on page 188
 - Configuring PoE (CLI Procedure) on page 193
 - Configuring PoE (J-Web Procedure) on page 195
 - PoE and EX Series Switches Overview on page 181

interval

Syntax	interval <i>minutes</i> ;
Hierarchy Level	[edit poe interface (all <i>interface-name</i>) telemetries]
Release Information	Statement introduced in JUNOS Release 9.0 for EX Series switches.
Description	Modify the interval for logging telemetries if you are monitoring the per-port power consumption for PoE interfaces.
Options	<i>minutes</i> —Frequency of logging. Range: 1 through 30 minutes Default: 5 minutes
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Topics	<ul style="list-style-type: none">■ Example: Configuring PoE Interfaces on an EX Series Switch on page 185■ Example: Configuring PoE Interfaces with Different Priorities on an EX Series Switch on page 188■ Configuring PoE (CLI Procedure) on page 193■ Configuring PoE (J-Web Procedure) on page 195■ PoE and EX Series Switches Overview on page 181

management

Syntax	management <i>type</i> ;
Hierarchy Level	[edit poe]
Release Information	Statement introduced in JUNOS Release 9.0 for EX Series switches. class option introduced in JUNOS Release 9.3 for EX Series switches.
Description	Designate the way that the switch's PoE controller allocates power to the PoE ports.
Default	static
Options	<p><i>type</i> —Management type:</p> <ul style="list-style-type: none"> ■ class—The power available for the interface is determined based class of powered device connected. See section Classes of Powered Devices in “PoE and EX Series Switches Overview” on page 181 for more information. ■ static—The switch reserves a certain amount of power for the PoE port even when a powered device is not connected to the port. This setting ensures that power is available when needed.
Required Privilege Level	<p>router—To view this statement in the configuration.</p> <p>router-control—To add this statement to the configuration.</p>
Related Topics	<ul style="list-style-type: none"> ■ Example: Configuring PoE Interfaces on an EX Series Switch on page 185 ■ Example: Configuring PoE Interfaces with Different Priorities on an EX Series Switch on page 188 ■ Configuring PoE (CLI Procedure) on page 193 ■ Configuring PoE (J-Web Procedure) on page 195 ■ PoE and EX Series Switches Overview on page 181

maximum-power

Syntax	maximum-power <i>watts</i> ;
Hierarchy Level	[edit poe interface (all <i>interface-name</i>)]
Release Information	Statement introduced in JUNOS Release 9.0 for EX Series switches.
Description	Maximum amount of power that can be supplied to the port.
Default	15.4 W
Options	<i>watts</i> Range: 0 through 15.4 Default: 15.4 W
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Topics	<ul style="list-style-type: none">■ Example: Configuring PoE Interfaces on an EX Series Switch on page 185■ Example: Configuring PoE Interfaces with Different Priorities on an EX Series Switch on page 188■ Configuring PoE (CLI Procedure) on page 193■ Configuring PoE (J-Web Procedure) on page 195■ PoE and EX Series Switches Overview on page 181

priority

Syntax	<code>priority value;</code>
Hierarchy Level	[edit poe interface (all <i>interface-name</i>)]
Release Information	Statement introduced in JUNOS Release 9.0 for EX Series switches.
Description	Set the priority for shutdown of individual ports when there is insufficient power for all PoE ports. If a port is set as high priority and a situation arises where there is not sufficient power for all the PoE ports, the available power is directed to the higher priority port(s). If the switch needs to shut down powered devices because a power supply fails and there is insufficient power, low priority devices are shut down before high priority devices.
Default	low
Options	<p><i>value</i> —high or low:</p> <ul style="list-style-type: none"> ■ high—Specifies that this port is to be treated as high priority in terms of power allocation. If there is insufficient power for all the PoE ports, the available power is directed to this port. If the switch needs to shut down powered devices because a power supply fails and there is insufficient power, the power is not shut down on this port until after it has been shut down on all the low priority ports. ■ low—Specifies that this port is to be treated as low priority in terms of power allocation. If there is insufficient power for all the PoE ports, power is not supplied to this port. If the switch needs to shut down powered devices because a power supply fails and there is insufficient power, the power is shut down on this port before it is shut down on high priority ports.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Topics	<ul style="list-style-type: none"> ■ Example: Configuring PoE Interfaces with Different Priorities on an EX Series Switch on page 188 ■ Configuring PoE (CLI Procedure) on page 193 ■ Configuring PoE (J-Web Procedure) on page 195 ■ PoE and EX Series Switches Overview on page 181

telemetries

Syntax `telemetries {
 disable;
 duration hours;
 interval minutes;
}`

Hierarchy Level [edit poe interface (all | *interface-name*)]

Release Information Statement introduced in JUNOS Release 9.0 for EX Series switches.

Description Allows you to log per port PoE power consumption.

If you want to log per-port power consumption, you must explicitly specify the **telemetries** statement. You can enable telemetries for all the PoE interfaces by setting `poe interface all`. However, if you modify the configuration of any individual PoE interface (for example, to change the **priority**, you must also specify the telemetries for that interface in order to maintain the logging. If you do not specify telemetries for a PoE interface, logging is disabled.

The statements are explained separately.

Default If the telemetries statement is specified, logging is enabled with the default values for **interval** and **duration**,

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

- Related Topics**
- Example: Configuring PoE Interfaces with Different Priorities on an EX Series Switch on page 188
 - Configuring PoE (CLI Procedure) on page 193
 - Configuring PoE (J-Web Procedure) on page 195
 - PoE and EX Series Switches Overview on page 181

Chapter 13

Operational Mode Commands for PoE

show poe controller

- Syntax** show poe controller
<detail | summary>
- Release Information** Command introduced in JUNOS Release 9.0 for EX Series switches.
- Description** Display the status of the Power over Ethernet (PoE) software module controller.
- Options** none—Display general parameters of the PoE software module controller.
detail | summary—(Optional) Display the specified level of output.
- Required Privilege Level** view
- Related Topics**
- show poe interface
 - Example: Configuring PoE Interfaces on an EX Series Switch on page 185
 - Example: Configuring PoE Interfaces with Different Priorities on an EX Series Switch on page 188
 - Monitoring PoE on page 197
- List of Sample Output** show poe controller on page 210
- Output Fields** Table 22 on page 210 lists the output fields for the show poe controller command. Output fields are listed in the approximate order in which they appear.

Table 22: show poe controller Output Fields

Field Name	Field Description
Ctrl-index	Identifies the controller.
Max-power	Specifies the maximum power that can be provided by the switch to PoE ports.
power-consumption	Specifies the total amount of power being used by the PoE ports, as measured by the specified telemetries settings.
Guard-band	Specifies the amount of power that has been placed in reserve.
Management	Specifies the management mode. Static is the only management mode supported.

```

show poe controller user@host> show poe controller

Ctrl-index  Max-power  power-consumption  Guard-band  Management
      0         305 W         0W             15W         Static
    
```

show poe interface

Syntax	show poe interface <ge-fpc/pic/port>
Release Information	Command introduced in JUNOS Release 9.0 for EX Series switches.
Description	Display the status of Power over Ethernet (PoE) ports.
Options	none—Display status of all PoE ports on the switch. ge-fpc/pic/port—(Optional) Display the status of a specific PoE port on the switch.
Required Privilege Level	view
Related Topics	<ul style="list-style-type: none"> ■ Example: Configuring PoE Interfaces on an EX Series Switch on page 185 ■ Example: Configuring PoE Interfaces with Different Priorities on an EX Series Switch on page 188 ■ Monitoring PoE on page 197
List of Sample Output	show status for all poe interfaces on the switch on page 211 show status for a specific PoE interface on the switch on page 212
Output Fields	Table 23 on page 211 lists the output fields for the show poe interface command. Output fields are listed in the approximate order in which they appear.

Table 23: show poe interface Output Fields

Field Name	Field Description
PoE Interface	Specifies the interface address.
Enabled	Specifies whether PoE capabilities are enabled or disabled.
status	Specifies whether PoE is currently being provided to the port.
max-power	Specifies the maximum power that can be provided to the port.
priority	Specifies whether the port is high or low priority.
power-consumption	Specifies how much power is being used by the port, as measured by the specified telemetries settings.
Class	Indicates the IEEE 802.af classification that defines the maximum power requirements for a powered device.

```

show status for all poe user@host> show poe interface
interfaces on the switch Interface Enabled status max-power priority power-consumption Class
ge-0/0/1 Enabled OFF 15.4W Low 0.0W 0
ge-0/0/3 Enabled OFF 12.0W High 0.0W 0
ge-0/0/5 Enabled OFF 15.4W Low 0.0W 0

```

```
show status for a user@host> show poe interface ge-0/0/3
specific PoE interface on
the switch      PoE interface status:
                PoE interface                : ge-0/0/3
                PoE capability of the interface : Enabled
                Current status of power supply on interface : OFF
                Power limit on the interface    : 12.0W
                Priority                        : High
                Power consumed                 : 0.0W
                Class of power device          : 0
```

show poe telemetries interface

Syntax	show poe telemetries interface <i>ge-fpc/pic/port</i> all x
Release Information	Command introduced in JUNOS Release 9.0 for EX Series switches.
Description	Display a history of power consumption on the specified interface.
Options	<p><i>ge-fpc/pic/port</i> —Display telemetries for the specified PoE interface.</p> <p>all—Display all telemetries records for the specified PoE interface.</p> <p>x —Display the specified number of telemetries records for the specified PoE interface.</p>
Required Privilege Level	view
Related Topics	<ul style="list-style-type: none"> ■ show poe interface ■ Example: Configuring PoE Interfaces on an EX Series Switch on page 185 ■ Example: Configuring PoE Interfaces with Different Priorities on an EX Series Switch on page 188 ■ Monitoring PoE on page 197
List of Sample Output	<p>show poe telemetries interface (Last 10 Records) on page 213</p> <p>show poe telemetries interface (All Records) on page 214</p>
Output Fields	Table 24 on page 213 lists the output fields for the show poe telemetries interface command. Output fields are listed in the approximate order in which they appear.

Table 24: show poe telemetries interface Output Fields

Field Name	Field Description
S1 No	Number of the record for the specified port. Record number 1 is the most recent.
Timestamp	Time that the power-consumption data was gathered.
Power	Amount of power provided by the specified port at the time the data was gathered.
Voltage	Maximum voltage provided by the specified port at the time the data was gathered.

```

show poe telemetries user@switch> show poe telemetries interface ge-0/0/0 10
interface ( Last 10
Records)
  S1 No   Timestamp                Power   Voltage
  1       01-27-2008 18:19:58 UTC    15.4W  51.6V
  2       01-27-2008 18:18:58 UTC    15.4W  51.6V
  3       01-27-2008 18:17:58 UTC    15.4W  51.6V
  4       01-27-2008 18:16:58 UTC    15.4W  51.6V

```

```

5      01-27-2008 18:15:58 UTC 15.4W  51.6V
6      01-27-2008 18:14:58 UTC 15.4W  51.6V
7      01-27-2008 18:13:58 UTC 15.4W  51.6V
8      01-27-2008 18:12:57 UTC 15.4W  51.6V
9      01-27-2008 18:11:57 UTC 15.4W  51.6V
10     01-27-2008 18:10:57 UTC 15.4W  51.6V

```

show poe telemetries interface (All Records) user@switch> **show poe telemetries interface ge-0/0/0 all**

Sl No	Timestamp	Power	Voltage
1	01-27-2008 18:19:58 UTC	15.4W	51.6V
2	01-27-2008 18:18:58 UTC	15.4W	51.6V
3	01-27-2008 18:17:58 UTC	15.4W	51.6V
4	01-27-2008 18:16:58 UTC	15.4W	51.6V
5	01-27-2008 18:15:58 UTC	15.4W	51.6V
6	01-27-2008 18:14:58 UTC	15.4W	51.6V
7	01-27-2008 18:13:58 UTC	15.4W	51.6V
8	01-27-2008 18:12:57 UTC	15.4W	51.6V
9	01-27-2008 18:11:57 UTC	15.4W	51.6V
10	01-27-2008 18:10:57 UTC	15.4W	51.6V
11	01-27-2008 18:09:57 UTC	15.4W	51.6V
12	01-27-2008 18:08:57 UTC	15.4W	51.6V
13	01-27-2008 18:07:57 UTC	15.4W	51.6V
14	01-27-2008 18:06:57 UTC	15.4W	51.6V
15	01-27-2008 18:05:57 UTC	15.4W	51.6V
16	01-27-2008 18:04:56 UTC	15.4W	51.6V
17	01-27-2008 18:03:56 UTC	15.4W	51.6V
18	01-27-2008 18:02:56 UTC	15.4W	51.6V
19	01-27-2008 18:01:56 UTC	15.4W	51.6V
20	01-27-2008 18:00:56 UTC	15.4W	51.6V
21	01-27-2008 17:59:56 UTC	15.4W	51.6V
22	01-27-2008 17:58:56 UTC	15.4W	51.6V
23	01-27-2008 17:57:56 UTC	15.4W	51.6V
24	01-27-2008 17:56:55 UTC	15.4W	51.6V
25	01-27-2008 17:55:55 UTC	15.4W	51.6V
26	01-27-2008 17:54:55 UTC	15.4W	51.6V
27	01-27-2008 17:53:55 UTC	15.4W	51.6V
28	01-27-2008 17:52:55 UTC	15.4W	51.6V
29	01-27-2008 17:51:55 UTC	15.4W	51.6V
30	01-27-2008 17:50:55 UTC	15.4W	51.6V
31	01-27-2008 17:49:55 UTC	15.4W	51.6V
32	01-27-2008 17:48:55 UTC	15.4W	51.6V
33	01-27-2008 17:47:54 UTC	15.4W	51.6V
34	01-27-2008 17:46:54 UTC	15.4W	51.6V
35	01-27-2008 17:45:54 UTC	15.4W	51.6V
36	01-27-2008 17:44:54 UTC	15.4W	51.6V
37	01-27-2008 17:43:54 UTC	15.4W	51.6V
38	01-27-2008 17:42:54 UTC	15.4W	51.6V
39	01-27-2008 17:41:54 UTC	15.4W	51.6V
40	01-27-2008 17:40:54 UTC	15.4W	51.6V
41	01-27-2008 17:39:53 UTC	15.4W	51.6V
42	01-27-2008 17:38:53 UTC	15.4W	51.6V
43	01-27-2008 17:37:53 UTC	15.4W	51.6V
44	01-27-2008 17:36:53 UTC	15.4W	51.6V